

The Wage Effects of Generative AI*

José Azar[†] Mireia Giné[‡] Javier Sanz-Espín[§]

August 21, 2026

Abstract

This paper studies the impact on wages and worker mobility of generative artificial intelligence. We use occupation-by-industry data from the Occupational Employment and Wage Statistics (OEWS) and data on postings and matched employer-employee data from Revelio. We find that moving from the 10th percentile to the 90th percentile of GenAI exposure is associated with a wage decline of 4.9% in OEWS data, with no statistically detectable change in employment. In Revelio data, moving from the 10th to the 90th percentile of exposure is associated with a decline of 8.71% in posted wages, and of 10.76% in starting wages. More exposed workers experience significantly lower job-to-job transition rates, but no change in transitions to persistent non-employment. Wage declines are largest among junior workers and occur across all age groups within junior positions. Finally, we find that AI exposure is associated with a decline in the elasticity of separation with respect to wages, which suggests that firms' labor market power has increased in occupations that are exposed to automation by GenAI. The combined evidence indicates that generative AI exposure has so far negatively affected workers' wages, mobility, and outside opportunities in the labor market.

Keywords: Generative AI; wages; firm organization; labor market power; monopsony; job ladders; occupational exposure.

JEL Codes: J23; J24; J31; J41; O33; D22.

*We thank seminar participants at IESE Business School, the European Central Bank, Tilburg University, Erasmus University Rotterdam, Bayes Business School, Queen Mary University, National University of Singapore, and the Institute for Employment Research (IAB) 2nd Workshop on Imperfect Competition in the Labor Market for helpful comments and suggestions. 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. This paper was previously circulated under the title 'AI Is Already Eroding Wages: Quasi-Experimental Evidence from Occupational Exposure'. All errors are our own.

[†]University of Navarra, School of Economics and Business, 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, School of Economics and Business. Email: javier.sanz-espín@tsm-education.fr.

1 Introduction

The public release of ChatGPT on November 30, 2022 triggered widespread concern that generative artificial intelligence (GenAI) would subject white-collar workers to a massive automation shock. Anthropic CEO Dario Amodei warned that GenAI could eliminate up to half of entry-level white-collar jobs.¹ However, the early empirical evidence points to limited aggregate employment effects so far. For example, [Hampole et al. \(2025\)](#) report small overall employment effects of GenAI exposure, and [Brynjolfsson et al. \(2025a\)](#) document a significant employment decline only among workers aged 22-25. Still, these findings do not necessarily imply an absence of labor market effects, as generative AI could be affecting wages and eroding workers' outside options even if firms are not (yet) eliminating white collar jobs *en masse*.

In this paper, we study the effect of GenAI on real wages and worker mobility in knowledge-intensive service industries. Using occupation-by-industry data from the Occupational Employment and Wage Statistics (OEWS) and data on job postings and matched employer-employee data from Revelio, we find that wages and separations declined significantly more in occupations that were more exposed to GenAI automation after the release of ChatGPT. At the same time, we find no significant negative effect on aggregate employment, and no increase in separations to persistent non-employment. The effects are identified by comparing the evolution of wages and worker mobility across occupations with different levels of exposure to GenAI automation within firms (or within industries in the case of the OEWS) before and after the release of ChatGPT. Thus, our results provide direct information about the evolution of wages in exposed occupations relative to less exposed occupations.

Quantitatively, we find that going from the 10th percentile of GenAI exposure to the 90th percentile is associated with a 4.9% decline in median hourly wages in the OEWS data in 2025, and with declines of 8.7% in posted wages and 10.8% in wages of new positions in Revelio by 2026.² The stronger results from Revelio are reasonable since they are based on wages for job postings and starting positions, which are expected to respond to labor market conditions faster than the level of wages for all active posi-

¹See [VandeHei and Allen \(2025\)](#).

²The Revelio data provides direct information on wages from job postings, and also combines the information from job postings with additional data on wages from self-reported salary information websites (such as `levels.fyi`), H-1B visa applications (LDA database), and the Current Population Survey to assign an imputed starting wage for each position in their dataset based on its characteristics (firm, job title, geography, worker seniority, and tenure at the firm). The latest Revelio vintage used in this paper covers data through June 2026.

tions. We also show that the relative decline in wages for higher exposed occupations to GenAI is not driven by changes in composition of workers in terms of average education, experience or demographic characteristics.

The literature on GenAI’s impact on employment has generally found stronger effects on younger and more junior workers ([Brynjolfsson et al., 2025a](#); [Hosseini and Lichtinger, 2025](#)). We analyze the heterogeneity in the effect of GenAI on wages by seniority and age. The age distributions overlap substantially across seniority groups, making it possible to compare workers of different ages holding positions at the same hierarchical level. The wage decline is highest for junior positions. By 2026, moving from the 10th to the 90th percentile of GenAI exposure is associated with wage declines of approximately 7.5% for junior positions and 3.2% for mid-level positions. For senior positions the estimate is not statistically significant. GenAI exposure is therefore associated with the largest wage decline at the lower rungs of the firm hierarchy. Within junior positions, we see significant wage declines across all age groups, including workers older than 40. Interestingly, mid-level position wage declines are concentrated among young workers, while there is no significant effect for senior workers in any age group. Thus, in contrast with the literature on the effects of GenAI on employment, we find that there are significant wage effects that go beyond the youngest and most junior workers.

Recent papers on the effects of GenAI on the labor market have highlighted two main threats to identification when using the 2022 ChatGPT shock: (i) correlation between AI exposure and the suitability of remote work ([Lambert and Schindler, 2026](#)), and (ii) heterogeneity in the sensitivity of different industries or firms to macroeconomic shocks, such as changes in interest rates ([Isenko and Curto Millet, 2026](#)). To address the first concern, we control for interactions of year fixed effects with occupational exposure to work-from-home (WFH), which we measure (following [Hansen et al., 2023](#)), based on the fraction of postings in an occupation that mention the possibility of remote work. In addition, to address the second concern, we control for the heterogeneity in sensitivity to macroeconomic shocks by including industry-year fixed effects in our OEWS specifications, and by including firm-year and MSA-year fixed effects in our Revelio specifications. These absorb any heterogeneous response to macroeconomic shocks at the industry level in the OEWS case, and at the firm or MSA level in the Revelio case. Furthermore, in robustness checks, we control for potential heterogeneity in the response of different occupations to macroeconomic shocks by including occupation-year fixed effects. All results are robust to

including these controls, which rules out WFH or the business cycle as the explanation for the observed pattern of wage declines.

Next, we explore whether workers in GenAI-exposed occupations have fewer outside opportunities post-ChatGPT. We start by estimating the effect of GenAI exposure on worker separations, as a proxy for quits. We distinguish between job-to-job transitions, defined as moves to another job within 90 days, and separations followed by persistent non-employment. Moving from the 10th to the 90th percentile of exposure is associated with a 3.0 percentage point decline in the separation rate. This is economically meaningful, as it is equivalent to 18.6% of the pre-ChatGPT mean in the separation rate. Most importantly, this decline comes from job-to-job transitions, which fall by 2.8 percentage points (28.8% of their pre-period mean). In contrast, we find no statistically significant change in separations to persistent non-employment. These results indicate that more exposed workers become less likely to leave their employers and, in particular, less likely to move directly to another job. Separation rates decline significantly across all seniority and age groups within the firm. The worker mobility effects are therefore more widespread across the firm hierarchy than the wage effects.

The negative effect of GenAI on separations suggests that exposed workers have worse outside options in the labor market. Previous work has found that, when outside opportunities deteriorate, separations become less responsive to wages, and conversely, when labor markets tighten, separations become more responsive (Depew and Sørensen, 2013; Hirsch et al., 2018; Autor et al., 2023). We explore whether this is also happening to workers exposed to the ChatGPT shock. In particular, we estimate workers' elasticity of separation with respect to posted wages. We use firms' posted wages as a measure of their wage policies, and instrument wages using an AKM-style (Abowd et al., 1999) firm wage component (as in Bassier et al., 2022). We find that the separation elasticity changes little after the introduction of ChatGPT for workers with low GenAI exposure, but declines significantly for workers that are highly exposed to GenAI automation.³ This evidence suggests that firms' wage-setting power has increased in occupations exposed to GenAI.

This paper's results speak mainly to three literatures. First, the paper's main contribution is to show that the generative AI shock has significantly reduced exposed workers' wages and separations (relative to less exposed workers), and has increased firms' labor market power over them. We also find that

³The average separation elasticity is -2.27 in the full sample, consistent with average elasticities found in meta-analyses of the monopsony literature (Sokolova and Sørensen, 2021; Caldwell et al., 2023).

the incidence of this adjustment is highly uneven within firms, with the largest wage declines occurring among junior workers at the base of the organization, across all age groups. These findings contribute to the literature on the labor market effects of automation and AI exposure ([Acemoglu and Restrepo, 2018, 2020](#); [Webb, 2020](#); [Hui et al., 2024](#); [Brynjolfsson et al., 2025a](#); [Hampole et al., 2025](#); [Hosseini and Lichtinger, 2025](#)). By showing relative wage declines in knowledge-intensive occupations, our results suggest that generative AI may be starting to drive a reversal of the classic skill-biased technological change (SBTC) patterns observed over recent decades ([Katz and Murphy, 1992](#); [Goldin and Katz, 2008](#); [Autor et al., 2020](#); [Katz, 2025](#)).

Second, we also contribute to the literature linking worker mobility to wage determination ([Burdett and Mortensen, 1998](#); [Postel-Vinay and Robin, 2002](#); [Karahan et al., 2017](#); [Autor et al., 2023](#); [Börschlein et al., 2024](#); [Bossler and Popp, 2026](#); [Heise et al., 2026](#)). The wage decline for GenAI exposed workers is accompanied by lower separation and job-to-job transition rates, but not by greater movement into persistent non-employment. These findings connect the wage response to a deterioration in workers' outside opportunities.

Finally, we contribute to the literature on monopsony power and its implications for firms' wage setting ([Card et al., 1993](#); [Card and Krueger, 1994](#); [Manning, 2003](#); [Ashenfelter et al., 2010](#); [Manning, 2011](#); [Sokolova and Sorensen, 2021](#); [Hirsch et al., 2022](#); [Azar et al., 2022](#); [Caldwell et al., 2023](#); [Azar and Marinescu, 2024a,b](#)). We show that the elasticity of separation with respect to wages declined more after ChatGPT for workers with higher levels of generative AI exposure. This suggests that GenAI has increased firms' monopsony power over workers in exposed occupations.

The remainder of the paper is organized as follows. Section 2 describes the data and the construction of generative AI exposure. Section 3 presents the empirical strategy. Section 4 reports the main results on wages, employment, worker mobility, posted-wage separation elasticities, and heterogeneity by seniority and age. Section 5 presents robustness checks. Section 6 concludes.

2 Data

We use two complementary sources of U.S. labor market data. The first is the Occupational Employment and Wage Statistics (OEWS) from the Bureau of Labor Statistics (BLS), which provides survey-based employment and wage information by detailed occupation and industry. The second is Revelio,

a large worker-level dataset constructed from online resumes and professional profiles, which allows us to study posted and starting wages of new positions and worker flows. It also allows us to estimate heterogeneous effects by worker age and seniority.

We focus on knowledge-intensive industries, corresponding to NAICS two-digit sectors 51-56: Information, Finance and Insurance, Real Estate, Professional, Scientific, and Technical Services, Management of Companies, and Administrative and Support and Waste Management Services. These industries rely heavily on cognitive, administrative, analytical, and information-processing tasks, making them a natural setting in which to study the early effects of large language models on the labor market.

Our baseline treatment variable is based on the occupation-level measure of generative AI exposure developed by [Eloundou et al. \(2024\)](#). The index measures the share of occupational tasks that large language models can perform or accelerate and ranges from zero to one. Because it is constructed from the pre-existing task content of occupations rather than post-ChatGPT labor market outcomes, it provides ex-ante variation in potential exposure to generative AI.

For the OEWS analysis, we assign exposure to the corresponding detailed occupation and hold it fixed over time. For the Revelio analysis, we begin with workers' detailed occupational codes and use the composition of each firm-by-six-digit-SOC-occupation-by-MSA cell on November 30, 2022 to construct exposure. The resulting measure, denoted by AI_{fom} , is the pre-ChatGPT employment-weighted average exposure of the detailed O*NET occupations represented in the cell. We hold this measure fixed throughout the sample period. Cells within the same six-digit SOC occupation can therefore have different exposure because their detailed 8-digit O*NET occupational composition differed before ChatGPT. This gives us predetermined variation in exposure across firms and local labor markets, even within the same six-digit occupation.

Figure 1 shows the employment-weighted distribution of GenAI exposure in the OEWS sample. The distribution is bimodal. A relatively small group of occupation-by-industry cells has exposure close to zero, while most of the employment mass is concentrated around 0.55. The interquartile range runs from 0.334 to 0.576. This dispersion motivates our continuous-treatment design, which uses variation in the intensity of exposure rather than dividing occupations into treated and untreated groups.

2.1 OEWS

The Occupational Employment and Wage Statistics (OEWS) from the Bureau of Labor Statistics (BLS) provides annual employment and wage information by detailed occupation and industry. We construct a balanced panel for 2018-2025 at the six-digit SOC occupation by four-digit NAICS industry by year level.⁴ Our wage outcome is the logarithm of the CPI-deflated median hourly wage. Our employment outcome is the logarithm of total employment in the corresponding occupation-by-industry cell. We retain cells observed in every year from 2018 through 2025 and use 2022 employment as the baseline employment weight. The final OEWS sample contains 24,736 occupation-by-industry-year observations, corresponding to 3,092 occupation-by-industry cells, 431 detailed occupations, and 37 four-digit industries. Table 1 reports summary statistics for this sample.

Because the sample spans multiple classification regimes, we harmonize both occupational and industry codes over time. Occupations are harmonized across the 2010 and 2018 SOC systems, while industries are harmonized across the 2012, 2017, and 2022 NAICS systems. The harmonization procedure constructs consistent occupation-by-industry cells over time and aggregates cells when classification changes make this necessary. Appendix B describes the procedure in detail. The resulting panel expresses all occupations in a common detailed SOC classification and all industries in a common four-digit NAICS classification. We merge generative AI and work-from-home exposure only after completing this harmonization, preventing classification changes from mechanically altering treatment status.

2.2 Revelio

Revelio provides worker-level employment histories with information on employers, occupations, metropolitan areas, job titles, employment spells, position seniority, worker age, and wages. The raw data for the US include more than 490 million observations, 148 million workers, and 10 million firms, with some employment histories extending back to 1950. We use all available employment-history records to construct annual employment stocks and worker flows at the firm-SOC6-MSA level. Our analysis covers 2018-2026. The 2026 data come from the latest Revelio vintage and extend through June

⁴In robustness exercises, we repeat the analysis after aggregating occupations to the SOC 3-digit level while retaining 4-digit NAICS industries. The estimates show that the results are not sensitive to the level of occupational aggregation.

of that year. As in OEWS, we restrict the sample to knowledge-intensive industries. We also require cells to be observed in every year of the sample and to have had more than one employee in 2022. The resulting balanced panel contains 15,218,901 cell-year observations and 1,690,989 firm-by-occupation-by-MSA cells. Although the panel is balanced in terms of active positions and separation rate measures, the wage panel is not balanced because some firm-occupation-MSA cells do not have new positions every year. However, we show in robustness checks that the results are robust to using a balanced wage panel (which represents around 14% of the observations in the balanced active positions sample).

We use two wage measures from Revelio. Our baseline outcome is the real log starting wage of newly created positions. Revelio estimates this measure by combining wage information from job postings, self-reported compensation websites, H-1B visa applications, and the Current Population Survey with information about the firm, job title, location, seniority, and worker tenure. We use starting wages as our main outcome because they draw on several complementary sources and provide considerably broader coverage than advertised wages alone.

Our second measure is the non-imputed wages reported directly in Revelio job postings. Posted wages provide an independently observed measure of firms' wage offers. We use them as an alternative outcome in the difference-in-differences analysis and as the measure of firms' wage policies in the separation-elasticity analysis. Posted wages are aggregated to the firm-occupation-MSA-year level, and the corresponding regressions are weighted by the cell's average number of vacancies during 2018-2022.

For the starting-wage analysis, we assign each newly created position to a year using its starting date.⁵ We retain positions with valid firm, occupation, and MSA identifiers and a positive starting wage. After deflating wages with the Consumer Price Index, we calculate the mean real log starting wage in each firm-occupation-MSA-year cell.

Revelio assigns positions to seven ordered seniority categories. We combine the first six into three broader groups. Positions with a seniority score of one or two are classified as junior, those with scores of three or four as mid-level, and those with scores of five or six as senior. We exclude the seventh category, which consists of top executives and is qualitatively different from the rest of the organizational hierarchy. We calculate wages for the full cell and separately by seniority. To distinguish hierarchical position from worker age, we also divide each seniority group into workers aged 22-30, 31-40, and 41-70.

⁵Positions beginning in December are assigned to the following year so that 2022 remains a fully untreated year following the public release of ChatGPT on November 30, 2022.

Worker flows are measured annually. We classify a separation as a job-to-job transition if the worker begins another position within 90 days. If no subsequent position is observed during that window, we classify it as a separation to persistent non-employment.

The separation, job-to-job transition, and persistent-non-employment rates divide the corresponding annual flow by lagged employment (measured as number of active positions) in the cell. Annual flows can exceed the initial employment stock in cells with substantial turnover during the year. To use a common and economically meaningful sample for all three outcomes, we exclude observations for which any of the rates exceeds one, as well as cells with missing or nonpositive lagged employment.

The separation-elasticity analysis uses the posted-wage sample matched to the employment histories of workers observed in the same firm, occupation, MSA, and year. Posted wages measure the wage policies workers face. For the benchmark markdown exercise, we estimate the separation elasticity as a continuous function of exposure and evaluate it at the mean exposure within each quintile.

Finally, we use the skills listed in workers' professional profiles to identify GenAI-related human capital. For every firm-by-occupation-by-MSA-year cell, we calculate the share of active workers who report at least one GenAI-related skill. We treat this share as a descriptive proxy for the diffusion of GenAI-related skills. It is not a direct measure of technology adoption, and our predetermined exposure measure remains the baseline treatment in the Revelio analysis.

We also classify newly created positions according to whether the worker had reported a GenAI-related skill by the year in which the position began. Starting wages can then be aggregated to the firm-by-occupation-by-MSA-by-year-by-GenAI-skill-status level. This allows us to compare workers with and without observed GenAI-related skills within the same labor market cell, while controlling for the overall share of GenAI-skilled workers in that cell. Because skills are self-reported and may themselves respond to wages or worker composition, we interpret both the worker indicator and the cell-level share as proxies for GenAI-related human capital rather than direct measures of AI use.

Table 2 presents weighted summary statistics for the full Revelio panel. Appendix Table A1 reports the corresponding statistics by seniority. Real log starting wages increase with seniority, from 10.007 among junior positions to 10.544 among mid-level positions and 11.035 among senior positions. As a final check on the wage data, Appendix Figure A1 compares average real log wages in Revelio and OEWS at the six-digit SOC occupation by four-digit NAICS industry level. The correlation is 0.81. This

is close to the correlation of 0.89 between Revelio job-based earnings and occupation-level earnings from the American Community Survey (ACS) reported by [Dorn et al. \(2025\)](#).

3 Empirical Strategy

Our empirical strategy uses the public release of ChatGPT on November 30, 2022 as a common shock. We compare how labor market outcomes changed after that date across occupations and labor market cells with different levels of predetermined GenAI exposure. Because the outcomes are measured annually, we treat 2022 as the final pre-ChatGPT year and use it as the reference year in the dynamic specifications. The post-ChatGPT period begins in 2023.

Exposure is continuous rather than binary. Our estimates therefore capture whether outcomes changed more in units with greater exposure; they do not compare a treated group with an untreated control group. In OEWS, treatment varies across detailed occupations. In Revelio, it varies across firm-by-six-digit-SOC-occupation-by-MSA cells according to their detailed O*NET occupational composition before ChatGPT. Appendix Figures [A2](#) and [A3](#) show that there is substantial variation in exposure both across broad occupational and industry groups and within them.

In the baseline specifications, we include both Generative AI and work-from-home exposure, both interacted with year fixed effects. GenAI exposure is defined at the occupation level in OEWS and at the firm-by-occupation-by-MSA level in Revelio. Work-from-home exposure is measured at the occupation level in both datasets. Including the two measures together helps distinguish the post-ChatGPT relationship with GenAI exposure from changes associated with occupations that were already more suited to remote work.

3.1 OEWS Specification

We first estimate the effect of generative AI exposure using the OEWS occupation-by-industry panel. Let o index detailed six-digit SOC occupations, j index four-digit NAICS industries, and t index years. Our dynamic specification is

$$y_{ojt} = \alpha_{oj} + \lambda_{jt} + \sum_{\tau \neq 2022} [\beta_{\tau} AI_o + \theta_{\tau} WFH_o] \mathbf{1}\{t = \tau\} + \varepsilon_{ojt}, \quad (3.1)$$

where y_{ojt} is either the real log median hourly wage or log employment. The variable AI_o denotes ex-ante generative AI exposure and WFH_o denotes pre-ChatGPT work-from-home exposure. Both measures are defined at the occupation level and remain fixed over time. The coefficients β_τ trace how wages or employment evolved in more exposed occupations relative to less exposed occupations, with 2022 as the omitted year. The work-from-home interactions allow occupations with different pre-existing remote-work suitability to follow different paths over time.

Occupation-by-industry fixed effects, α_{oj} , absorb permanent differences across cells, including differences in wage levels, employment, task content, and the way occupations are distributed across industries. Industry-by-year fixed effects, λ_{jt} , absorb shocks affecting all occupations in the same four-digit industry and year. The identifying comparison is therefore between occupations with different GenAI exposure operating in the same industry and year.

We also estimate a pooled specification in which the year-specific interactions are replaced by $AI_o \times Post_t$ and $WFH_o \times Post_t$, where $Post_t$ equals one from 2023 onward. The pooled coefficient summarizes the average post-ChatGPT differential, while the dynamic coefficients show how that differential evolves and allow us to examine pre-trends. Because both outcomes are in logarithms, the coefficients are semi-elasticities. A change in exposure of ΔAI implies an approximate outcome change of $100\beta_\tau \Delta AI$ percent.

Occupations with high GenAI exposure may differ systematically from occupations with low exposure. To improve comparability across the exposure distribution, we use stabilized inverse-probability weights for a continuous treatment. Let $i = (o, j)$ denote an occupation-by-industry cell, and let X_i include pre-ChatGPT work-from-home exposure and broad occupation and industry characteristics. We construct

$$\omega_i = \frac{\hat{f}(AI_i)}{\hat{f}(AI_i | X_i)}, \quad (3.2)$$

where the numerator is the estimated marginal density of GenAI exposure and the denominator is its conditional density given X_i . In the baseline procedure, we estimate both densities nonparametrically after residualizing exposure with respect to the conditioning variables. The final regression weight is the product of ω_i and the cell's employment in 2022. Standard errors are clustered by three-digit NAICS industry.

3.2 Revelio Specification

Using Revelio, we compare cells with different levels of predetermined generative AI exposure within the same firm and year while also accounting for shocks specific to local labor markets. The unit of observation is a firm-by-six-digit-SOC-occupation-by-MSA-by-year cell. Treatment, AI_{fom} , is constructed using the cell's detailed O*NET occupational composition on November 30, 2022 and is held fixed thereafter. Exposure can therefore differ across firms and local labor markets even within the same six-digit SOC occupation.

Let f index firms, o occupations, m metropolitan areas, and t years. For outcome k , we estimate

$$y_{fomt}^k = \alpha_{fom} + \lambda_{ft} + \mu_{mt} + \sum_{\tau \neq 2022} \left[\beta_{\tau}^k AI_{fom} + \theta_{\tau}^k WFH_o \right] \mathbf{1}\{t = \tau\} + \varepsilon_{fomt}^k. \quad (3.3)$$

Firm-by-occupation-by-MSA fixed effects, α_{fom} , absorb permanent differences across labor market cells. Firm-by-year fixed effects, λ_{ft} , absorb shocks affecting the whole firm in a given year, such as changes in product demand, financing conditions, hiring policies, or compensation practices. MSA-by-year fixed effects, μ_{mt} , absorb changes in local labor market conditions.

The key comparison is therefore between occupations with different predetermined exposure inside the same firm and year, after removing shocks common to the local labor market. Identification comes both from differences across six-digit occupations and from differences in detailed O*NET composition within a six-digit occupation.

We test the sensitivity of this comparison using more demanding fixed effects. We first add six-digit occupation-by-year fixed effects, which absorb national shocks affecting a detailed occupation in a given year. We then include occupation-by-MSA-by-year fixed effects, which absorb shocks to that occupation in a particular local labor market. Our most saturated specification includes both firm-by-occupation-by-year and occupation-by-MSA-by-year fixed effects. In this case, identification relies on predetermined differences in detailed O*NET composition across firms and MSAs within the same six-digit occupation.

As in OEWS, we estimate both dynamic and pooled specifications. The pooled model replaces the year-specific interactions with $AI_{fom} \times Post_t$ and $WFH_o \times Post_t$. Both specifications use the full 2018-2026 panel. Standard errors are two-way clustered by firm and occupation.

For the wage regressions, y_{fomt}^k is the mean real log starting wage of newly created positions. For the worker-flow regressions, the outcomes are the separation rate, job-to-job transition rate, and persistent-non-employment separation rate. Each rate divides the corresponding annual flow by lagged cell employment. We exclude cells with missing or nonpositive lagged employment and observations for which any of the three rates exceeds one. This keeps the estimation sample common across the worker-flow outcomes. The interpretation of β_τ^k depends on the dependent variable. For log wages, it is a semi-elasticity. For worker flows, it measures the change in the corresponding rate associated with a one-unit increase in exposure.

All weights are fixed using pre-ChatGPT information. Wage regressions are weighted by the average number of newly created positions in the cell during 2018-2022. Worker-flow regressions use average cell employment over the same period. The seniority-specific wage regressions use the average number of newly created positions in the relevant seniority group during 2018-2022. Throughout the paper, unless otherwise indicated, the 10th and 90th percentiles of exposure are calculated in the corresponding estimation sample using the regression weights. For log outcomes, we report the exact percentage change, $100[\exp\{\beta(P_{90} - P_{10})\} - 1]$. For worker-flow rates, which are expressed between zero and one, we report $100\beta(P_{90} - P_{10})$ percentage points. We also express this change relative to the weighted pre-ChatGPT mean of the corresponding outcome as $100\beta(P_{90} - P_{10})/\bar{y}_{pre}$ percent.

We construct occupational work-from-home exposure following [Hansen et al. \(2023\)](#). For each detailed occupation, we calculate the share of Revelio job postings classified as remote during 2021 and the pre-ChatGPT months of 2022. We average that share over the pre-period and hold it fixed thereafter. The measure captures an occupation’s pre-existing suitability for remote work rather than changes in remote-work adoption after ChatGPT. We interact it with year indicators in the dynamic regressions and with the post-ChatGPT indicator in the pooled regressions.

We also re-estimate the wage specifications using advertised wages in place of Revelio’s modeled starting wages. These regressions include the same firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects, as well as the same work-from-home interactions. They are weighted by the average number of vacancies in the cell during 2018-2022.

Finally, we use the GenAI skills reported in workers’ profiles in a set of complementary exercises. We first ask whether predetermined exposure predicts the subsequent diffusion of the GenAI worker share.

Appendix Figure A4 reports this event study. Appendix Table A2 then compares the wage relationship across cells that ever report a positive GenAI worker share and cells that do not. We also compare the starting wages of workers with and without reported GenAI-related skills, controlling alternatively for predetermined exposure, the cell-level GenAI worker share, or both.

These skill-based exercises are descriptive. Reported skills and their prevalence may respond to wages, worker composition, or other changes within the cell. We therefore do not interpret them as causal effects of GenAI adoption or skill acquisition. The baseline estimates remain reduced-form estimates based on predetermined cell-level exposure.

3.3 Posted-Wage Separation Elasticities

The worker-flow results show that exposed workers become less likely to leave their employers and, in particular, less likely to move directly to another job. We next ask whether they also become less responsive to the wages offered by their current employer. This elasticity exercise builds on the more direct evidence on worker mobility; it does not identify a separate mechanism.

Posted wages may respond to expected turnover or local labor market conditions. We therefore instrument log posted wages using an AKM-style firm wage component. Following the cross-fitting approach in Bassier et al. (2022), we randomly divide occupations into two non-overlapping groups, $\mathcal{G}_1$ and $\mathcal{G}_2$, and estimate the following posted-wage decomposition separately in each group:

$$\log w_{ifomt}^p = \alpha_i + \psi_f + \gamma_o + \rho_m + \kappa_t + u_{ifomt}, \quad (3.4)$$

where α_i , ψ_f , γ_o , ρ_m , and κ_t are worker, ultimate-parent firm, occupation, MSA, and year effects. For workers in occupations belonging to $\mathcal{G}_1$, we use the firm component estimated from $\mathcal{G}_2$ as the instrument. For workers in $\mathcal{G}_2$, we use the component estimated from $\mathcal{G}_1$. The instrument is therefore constructed using wage information from other occupations rather than the worker's own occupational group.

We begin with a linear specification. It provides a simple summary of how the relationship between wages and separations varies with GenAI exposure and changes after ChatGPT. We center exposure so that the main wage coefficients have a natural interpretation. Let

$$\widetilde{AI}_o = AI_o - \overline{AI}, \quad (3.5)$$

where $\bar{AI}$ is the exposure value used as the centering point. For the overall regression, this is the mean exposure in the full IV sample. For the seniority regressions, we use a common mean calculated across the three seniority samples. We then estimate

$$\begin{aligned}
S_{ifomt} = & \alpha_i + \gamma_{mo} + \lambda_{mt} + \delta_{ot} \\
& + \pi_0 \log w_{fomt}^P + \pi_1 \widetilde{AI}_o \log w_{fomt}^P \\
& + \pi_2 \log w_{fomt}^P \times \text{Post}_t + \pi_3 \widetilde{AI}_o \log w_{fomt}^P \times \text{Post}_t + \varepsilon_{ifomt}.
\end{aligned} \tag{3.6}$$

The dependent variable, S_{ifomt} , equals one when worker i separates from the firm. Worker fixed effects, α_i , absorb permanent differences in workers' separation propensities. MSA-by-occupation fixed effects, γ_{mo} , absorb permanent differences across local occupational labor markets. MSA-by-year fixed effects, λ_{mt} , account for changing local labor market conditions, while occupation-by-year fixed effects, δ_{ot} , absorb national shocks affecting a particular occupation. Standard errors are clustered by MSA-by-occupation labor market.

The coefficient π_0 measures the effect of log posted wages on the separation probability at the exposure centering point before ChatGPT. The coefficient π_1 shows how this wage response varies with exposure in the pre-period. The coefficient π_2 captures the post-ChatGPT change at the centering point, and π_3 shows whether that change becomes larger or smaller as exposure rises. Because the wage slopes are negative, positive values of π_2 and π_3 mean that separations become less responsive to wages after ChatGPT and that this attenuation grows with exposure. The main effects of exposure, the post indicator, and their interaction are absorbed by the occupation-by-year fixed effects.

All terms containing $\log w_{fomt}^P$ are treated as endogenous. We instrument $\log w_{fomt}^P$, $\widetilde{AI}_o \log w_{fomt}^P$, $\log w_{fomt}^P \times \text{Post}_t$, and $\widetilde{AI}_o \log w_{fomt}^P \times \text{Post}_t$ with $\widehat{\psi}_f^{CF}$, $\widetilde{AI}_o \widehat{\psi}_f^{CF}$, $\widehat{\psi}_f^{CF} \times \text{Post}_t$, and $\widetilde{AI}_o \widehat{\psi}_f^{CF} \times \text{Post}_t$, respectively.

To summarize the overall magnitude of the separation response, we evaluate the wage slope at the sample means of centered exposure, the post indicator, and their interaction. We divide this slope by the average fitted separation probability from a separate logit model estimated on the same IV sample:

$$\bar{\eta}^S = \frac{\hat{\pi}_0 + \hat{\pi}_1 \widetilde{AI} + \hat{\pi}_2 \overline{\text{Post}} + \hat{\pi}_3 \widetilde{AI} \times \overline{\text{Post}}}{\bar{\hat{s}}}, \quad (3.7)$$

where $\bar{\hat{s}}$ is the average fitted separation probability over the full estimation period. This statistic is the ratio of the average wage slope to the average fitted separation probability. It is not the average of worker-specific elasticities. Its standard error reflects uncertainty in the IV wage-slope estimates but treats the fitted separation probability as fixed.

We estimate this model for the full sample and separately for junior, mid-level, and senior workers. Seniority is fixed at its 2022 value so that changes in a worker's position after ChatGPT cannot determine the group to which the worker belongs. We exclude workers without a valid 2022 seniority classification and those in Revelio's executive category. Each seniority-specific elasticity uses the sample means and fitted separation probability from the relevant estimation sample.

The linear model is useful for summarizing the average relationship, but it requires the exposure gradient to be linear both before and after ChatGPT. We therefore also estimate a more flexible specification. We use five common knots to construct four restricted-cubic-spline basis functions, $B_r(AI_o)$ for $r = 1, \dots, 4$, and define $B_0(AI_o) = 1$. We estimate

$$S_{ifomt} = \alpha_i + \gamma_{mo} + \lambda_{mt} + \delta_{ot} + \sum_{r=0}^4 \left[\beta_{r0} B_r(AI_o) \log w_{fomt}^P + \beta_{r1} B_r(AI_o) \log w_{fomt}^P \times \text{Post}_t \right] + \varepsilon_{ifomt}. \quad (3.8)$$

The spline uses the same worker, MSA-by-occupation, MSA-by-year, and occupation-by-year fixed effects as the linear model. Its basis is constructed once from the pooled estimation sample, so the same functional form is used before and after ChatGPT.

As in the linear model, every term containing posted wages is endogenous. For each $r = 0, \dots, 4$, we instrument $B_r(AI_o) \log w_{fomt}^P$ and $B_r(AI_o) \log w_{fomt}^P \times \text{Post}_t$ with $B_r(AI_o) \hat{\psi}_f^{CF}$ and $B_r(AI_o) \hat{\psi}_f^{CF} \times \text{Post}_t$, respectively, where $\hat{\psi}_f^{CF}$ is the cross-fitted AKM firm wage component estimated using the other occupation group.

We convert the estimated wage slopes into elasticities using fitted separation probabilities from a separate logit model estimated on the same sample. The logit contains the same spline basis, the post-

ChatGPT indicator, and their interactions. Let $\hat{b}_p(x)$ denote the estimated effect of log posted wages on the separation probability at exposure x in period p , and let $\hat{s}_p(x)$ denote the fitted separation probability. The elasticity is

$$\hat{\eta}_p^S(x) = \frac{\hat{b}_p(x)}{\hat{s}_p(x)}. \quad (3.9)$$

Figures 7 and 9 evaluate this continuously estimated function at the mean exposure within each decile of the relevant estimation sample. The deciles are therefore display points rather than separate treatment groups. The confidence intervals reflect uncertainty in the IV wage-slope estimates while treating the fitted separation probabilities as fixed.

3.3.1 Benchmark Mapping from Separation Elasticities to Wage Markdowns

We use the estimated separation elasticities to calculate the wage markdowns implied by the canonical inverse-elasticity benchmark. This is an illustrative calibration, not a direct estimate of firms' markdowns or workers' productivity.

Following Manning (2003), we approximate the firm-level labor-supply elasticity as

$$\hat{\varepsilon}_p^L(x) = 2 \left| \hat{\eta}_p^S(x) \right|, \quad (3.10)$$

where $p \in \{\text{pre}, \text{post}\}$ denotes the period and $x \in [0, 1]$ denotes generative AI exposure. This approximation imposes the steady-state restriction that recruitment and separation elasticities are equal in absolute value. Under the inverse-elasticity rule, the implied relative wage markdown is

$$\hat{v}_p(x) \equiv \frac{\widehat{MRPL}_p(x) - \bar{w}_p(x)}{\bar{w}_p(x)} = \frac{1}{\hat{\varepsilon}_p^L(x)} = \frac{1}{2 \left| \hat{\eta}_p^S(x) \right|}. \quad (3.11)$$

This mapping assumes that firms face the residual labor-supply curve as the relevant constraint on wage setting and that wages satisfy the inverse-elasticity rule. It may fail in the presence of bargaining, non-wage amenities, efficiency wages, differences in worker productivity or labor-supply elasticities, internal pay constraints, or norms of horizontal equity (Caldwell et al., 2023). We therefore treat the resulting markdowns as an indicator of firms' power to set wages below the marginal revenue product of labor, rather than as direct estimates of the gap between the two.

3.4 Identification

The reduced-form designs require outcomes in units with different exposure to have followed parallel trends in the absence of ChatGPT, conditional on the fixed effects, work-from-home controls, and weighting procedures. The pre-2022 event-study coefficients provide a direct way to assess this assumption.

Because ChatGPT became publicly available at the same time for everyone, the design does not face the complications associated with staggered treatment adoption. But the common timing also limits what we identify. Exposure measures susceptibility to the technology rather than actual adoption, and even occupations with low direct exposure may be affected through spillovers or general-equilibrium effects. Our coefficients therefore capture differential changes across the exposure distribution, not the aggregate effect of ChatGPT on the labor market.

The main threat is that other post-2022 shocks may also have affected highly exposed occupations. Professional and white-collar occupations, for example, may have responded differently to the interest-rate cycle or the post-pandemic adjustment ([Ischenko and Curto Millet, 2026](#)). In OEWS, four-digit industry-by-year fixed effects absorb shocks common to all occupations in the same industry. The inverse-probability weights and work-from-home controls further adjust for differences related to broad occupational characteristics and remote-work suitability.

Firm-by-year fixed effects absorb every shock common to a firm in a given year, while MSA-by-year fixed effects absorb changes in local labor market conditions. The robustness specifications go further by controlling for six-digit occupation-by-year, occupation-by-MSA-by-year, and firm-by-occupation-by-year fixed effects. These absorb national occupation-specific shocks, shocks to an occupation in a particular local labor market, and shocks to a particular occupation within a firm. In the most saturated model, identification comes from predetermined differences in detailed O*NET composition across firms and MSAs within the same six-digit occupation.

These fixed effects substantially narrow the set of alternative explanations, but they do not eliminate it. A causal interpretation still requires any remaining post-2022 shocks to be unrelated to predetermined GenAI exposure after conditioning on the included controls.

A related concern is that GenAI exposure is correlated with occupations' suitability for remote work ([Lambert and Schindler, 2026](#)). We address this directly by including pre-ChatGPT work-from-home

exposure in both the pooled and dynamic regressions. In OEWS, work-from-home exposure also enters the model used to construct the inverse-probability weights. The robustness analysis reports estimates without this control, under alternative weighting procedures and exposure measures, in a balanced wage panel, using posted rather than modeled starting wages, and under alternative fixed-effect structures. We also use the prevalence of GenAI-related skills as complementary descriptive evidence, while keeping predetermined exposure as the treatment in the baseline design.

4 Results

4.1 Wages and Employment in OEWS

We begin with the OEWS evidence on wages and employment. The expected effect of GenAI exposure is not clear a priori. Automation can reduce demand for the tasks performed by exposed workers, but productivity gains and the creation of new tasks may offset some of that displacement ([Acemoglu and Restrepo, 2018](#); [Hui et al., 2024](#); [Hampole et al., 2025](#)).

Figure 2 presents the wage event study, Table 3 reports the corresponding pooled and dynamic estimates. There is no systematic differential trend before 2022. The coefficient turns negative in 2023, although it is not statistically significant, and becomes significant in 2024 and 2025. The estimates are -0.114 in 2024 and -0.083 in 2025. Moving from the 10th to the 90th percentile of exposure, these coefficients imply wage declines of approximately 6.7% and 4.9%, respectively. The pooled coefficient is -0.078 , corresponding to an average decline of approximately 4.6% over the same exposure contrast.

No comparable pattern appears in employment. Figure 3 shows that the post-ChatGPT coefficients are statistically insignificant and do not follow a sustained downward path. The pooled employment estimate is also insignificant. Thus, more exposed occupation-by-industry cells experience lower wages without a detectable decline in employment during the period covered by OEWS.

The regressions include GenAI and work-from-home exposure jointly. As Panel B of Table 3 shows, the negative wage estimates cannot be explained by occupations' pre-existing suitability for remote work. The contrast between wages and employment is therefore the first important result of the paper. The early adjustment to GenAI exposure appears in the price of labor before it is visible in aggregate headcount.

4.2 Wages and Worker Flows at the Firm Level

We next turn to Revelio, which allows us to compare occupations with different predetermined exposure within the same firm and year. Figure 4 and the first two columns of Table 4 report the results for the starting wages of newly created positions. The pre-ChatGPT coefficients are close to zero, while every post-ChatGPT coefficient is negative and statistically significant. The pooled coefficient is -0.140 . Moving from the 10th to the 90th percentile of exposure, this estimate implies an average starting-wage decline of approximately 5.8%. The coefficient becomes progressively more negative and reaches -0.266 in 2026, corresponding to a 10.8% decline over the same exposure contrast. The 2026 estimate should be interpreted with some caution because the latest Revelio vintage ends in June 2026.

Changes in the observable composition of new workers explain little of this decline. Table 5 adds controls for gender and race, college education, and experience. The baseline coefficient remains -0.140 after controlling for demographic composition, changes to -0.142 when college education is added, and falls only modestly to -0.132 after controlling for experience.

The effect is not uniform across industries or firms. Table 6 reports the largest coefficient in Information, at -0.195 , followed by Finance and Insurance, at -0.127 , and Professional Services, at -0.092 . The estimates for Real Estate, Management, and Administrative Services are smaller and statistically insignificant. The wage decline also increases with firm size. The coefficient is -0.066 among firms with fewer than 50 employees, -0.082 among firms with 50 to 250 employees, and -0.146 among firms with more than 250 employees.

We also study how worker flows changed as wages declined. Figure 5 reports the event study for the separation rate. The coefficients become negative immediately after 2022 and remain negative throughout the post-period. The pooled coefficient in Table 4 is -0.074 . Moving from the 10th to the 90th percentile of exposure, this estimate implies a decline of 3.0 percentage points in the separation rate, equivalent to 18.6% of its pre-ChatGPT mean.

Figure 6 separates job-to-job transitions from movements into persistent non-employment. Panel A shows a pronounced decline in job-to-job transitions. The pooled coefficient is -0.069 , implying a reduction of 2.8 percentage points between the 10th and 90th percentiles of exposure, or 28.8% of the pre-ChatGPT mean. Job-to-job transitions therefore account for approximately 92% of the implied decline in total separations. Panel B shows no comparable change in persistent non-employment. The

estimated decline is only 0.2 percentage points, or 3.9% of its pre-ChatGPT mean, and is not statistically significant.

These results point to a decline in labor market dynamism rather than widespread displacement. Workers in more exposed cells become less likely to leave their employer and, in particular, less likely to move directly to another job. They do not become more likely to enter persistent non-employment. This distinction matters because employer-to-employer mobility is an important source of wage growth and wage pressure (Burdett and Mortensen, 1998; Postel-Vinay and Robin, 2002; Karahan et al., 2017; Autor et al., 2023; Heise et al., 2026). Fewer outside opportunities can reduce wages even when existing employment relationships remain in place. This helps explain why we observe sizable wage effects without a corresponding decline in aggregate employment.

4.3 Posted-Wage Separation Elasticities

The wage and mobility results all point in the same direction: exposed workers face fewer attractive outside opportunities after ChatGPT. We now ask whether this change is also visible in the responsiveness of separations to the wages firms offer. We estimate the linear and spline AKM-IV models described in Section 3.3, instrumenting posted wages and their interactions with the cross-fitted AKM firm wage component.

4.3.1 Estimated separation elasticities

Table 8 first reports the linear specification. In the full sample, the coefficient on log posted wages is -0.188 . Higher posted wages therefore reduce the probability of separation at mean exposure before ChatGPT. The interaction with centered GenAI exposure is -0.415 , indicating that this negative relationship becomes stronger as exposure rises in the pre-period.

Both post-ChatGPT interactions have the opposite sign. The interaction between log posted wages and the post indicator is 0.096 , while the triple interaction with GenAI exposure is 0.260 . The negative wage slope therefore moves toward zero after ChatGPT, and the attenuation becomes larger as exposure rises. All four coefficients are significant at the one-percent level.

The same qualitative pattern appears at every seniority level. Before ChatGPT, the coefficients on posted wages and their interactions with exposure are negative for junior, mid-level, and senior workers.

The post-ChatGPT interactions are positive and statistically significant in all three groups. The decline in separation responsiveness is therefore not confined to one part of the firm hierarchy. The average separation elasticity is -2.272 in the full sample. It is -2.528 for junior workers, -3.661 for mid-level workers, and -3.228 for senior workers. On average, junior workers are the least responsive to posted wages, while mid-level workers are the most responsive.

This ranking is relevant because junior workers also experience the largest decline in starting wages. Lower wage responsiveness and larger wage declines are consistent with weaker mobility-based discipline on junior wages. They do not, however, establish that the difference in elasticities causes the seniority gradient in wages. The overall elasticity also need not equal a weighted average of the three seniority estimates. The seniority samples exclude workers without a valid 2022 classification and workers in Revelio’s executive category, and each model is estimated separately. The first stages are strong. The Kleibergen-Paap F -statistic is 194.03 in the full sample, 184.90 for junior workers, 229.25 for mid-level workers, and 67.64 for senior workers.

The linear estimates provide a useful summary, but they impose a linear relationship between GenAI exposure and separation responsiveness. Figure 7 relaxes this restriction. The spline specification also has a strong joint first stage, with a Kleibergen-Paap F -statistic of 88.7. The estimated elasticities are negative before and after ChatGPT, confirming that higher posted wages reduce separations. Before ChatGPT, their absolute magnitude rises with GenAI exposure and does so particularly sharply in the upper tail. After ChatGPT, the elasticities are closer to zero throughout the displayed distribution. The gap between the two periods becomes much larger in the highest exposure deciles.

Figure 8 plots this gap directly. The post-minus-pre change in the absolute elasticity, $\Delta|\eta^S|(x) = |\eta_{\text{post}}^S(x)| - |\eta_{\text{pre}}^S(x)|$ is negative at every displayed decile and becomes markedly more negative toward the top of the exposure distribution. The flexible estimates therefore confirm the message from the linear model: worker separations become less responsive to posted wages after ChatGPT, especially in highly exposed occupations.

Figure 9 shows that the high-exposure pattern appears within each seniority group. At lower deciles, the pre-post differences are smaller and vary somewhat across groups. Near the top of the distribution, the post-ChatGPT elasticities move substantially closer to zero for junior, mid-level, and senior workers alike. The change in separation responsiveness is therefore broader than the seniority gradient in wages.

This evidence is consistent with earlier work showing that separation and labor-supply elasticities fall when outside opportunities become scarcer (Depew and Sørensen, 2013; Hirsch et al., 2018). Conversely, the post-pandemic tightening increased the wage responsiveness of job-to-job separations among workers toward the bottom of the wage distribution (Autor et al., 2023). In our setting, the decline in responsiveness is greatest at high levels of GenAI exposure, where job-to-job mobility also falls. Worker mobility consequently places less pressure on the wages offered by firms.

This interpretation does not require an independent structural increase in monopsony power. The same pattern could arise if GenAI reduces firms’ relative demand for exposed tasks, if employers become more selective when filling exposed positions, or if workers find it difficult to move into less exposed work. Our estimates do not distinguish among these channels. What they share is a deterioration in workers’ outside opportunities and, with it, greater scope for firms to set wages without inducing departures.

4.3.2 Benchmark implications for wage markdowns

Table 9 applies the inverse-elasticity benchmark described in Section 3.3.1. We evaluate the continuously estimated spline at the mean exposure within each quintile and use the resulting separation elasticities to calculate implied firm-level labor-supply elasticities and relative wage markdowns. The quintiles are evaluation points, not separate treatment groups.

Before ChatGPT, the absolute separation elasticity increases from 1.874 at the first quintile mean to 5.158 at the fifth. After ChatGPT, it is lower at all five points and ranges from 1.318 to 2.947. The corresponding firm-level labor-supply elasticities range from 3.748 to 10.316 before ChatGPT and from 2.635 to 5.893 afterward. For the first four quintile means, the post-ChatGPT estimates are close to the average firm-level wage elasticity of approximately three reported by Caldwell et al. (2023). This comparison provides context for their magnitude but does not validate the assumptions behind the mapping.

The implied relative markdown, $\nu = (MRPL - w)/w$, rises at every evaluation point. The increase ranges from 38.2% at the first quintile mean to 75.1% at the fifth. Under the canonical inverse-elasticity rule, the decline in separation responsiveness therefore maps into a wider wage-setting wedge, with the largest proportional increase near the top of the exposure distribution. These calculations are cali-

brations rather than separate estimates of firms' structural markdowns. They show what the estimated elasticities would imply under the inverse-elasticity benchmark; they do not provide independent evidence about marginal products or worker productivity.

4.4 Wage Effects by Seniority

We next examine whether the wage response differs across the firm hierarchy. This distinction is important because GenAI may substitute for entry-level tasks while also making less experienced workers more productive (Noy and Zhang, 2023; Brynjolfsson et al., 2025b,a; Hosseini and Lichtinger, 2025).

Figure 10 and Table 7 show a clear seniority gradient. The pooled coefficient is -0.100 for junior positions, -0.060 for mid-level positions, and -0.043 for senior positions. The junior and mid-level estimates are statistically significant; the senior estimate is not. To make the percentage changes comparable across seniority groups, we hold fixed the 10th and 90th percentiles from the baseline weighted Revelio wage sample. Moving from the 10th to the 90th percentile of exposure, the pooled coefficients imply average wage declines of approximately 4.2% for junior positions, 2.5% for mid-level positions, and 1.8% for senior positions.

The dynamic estimates follow the same ordering. Junior wages fall in every post-ChatGPT year, with the coefficient reaching -0.181 in 2026. The mid-level effect is smaller but remains negative and statistically significant, reaching -0.075 in 2026. The senior coefficient is negative and significant in 2023 but becomes smaller and less precise afterward.

By 2026, moving from the 10th to the 90th percentile of exposure is associated with wage declines of approximately 7.5% for junior positions, 3.2% for mid-level positions, and 1.8% for senior positions. The senior estimate is not statistically significant. The wage effect is therefore not limited to junior work, but it is clearly largest and most persistent near the bottom of the hierarchy. Panel B of Table 7 shows no comparable seniority gradient in the work-from-home coefficients. The stronger junior response is therefore unlikely to be explained by junior positions being differentially sensitive to remote work.

4.5 Wage and Separation Effects by Age and Seniority

Age and seniority are related, but they do not measure the same thing. Figure 11 shows substantial overlap in the age distributions of junior, mid-level, and senior workers. Many older workers remain in

junior positions, while some younger workers occupy positions higher in the firm hierarchy. This allows us to distinguish the effect of rank from the effect of age.

Panel A of Figure 12 reports wage estimates for each seniority-by-age group. The coefficients are negative for junior workers aged 22-30, 31-40, and 41-70. The largest point estimate is, in fact, among junior workers older than 40. Among mid-level workers, the negative wage response is concentrated among those below age 40. The estimate for older mid-level workers is smaller and statistically insignificant. The senior estimates are also smaller and imprecisely estimated across all three age groups. The large junior effect is therefore not simply a young-worker effect. Wage declines appear throughout the junior age distribution and are largest among older juniors.

Panel B reports the results for separation rates. Here the pattern is broader. The estimates are negative and statistically significant in all nine seniority-by-age groups. The largest point estimates occur among junior workers aged 31-40 and mid-level workers below age 40, but separations also decline among older mid-level workers and among senior workers of every age group.

Wages and mobility therefore follow different seniority patterns. The wage decline is concentrated toward the bottom of the hierarchy, while the decline in separations extends across seniority levels and age groups. The particularly large wage response among older junior workers could reflect differences in outside opportunities or greater difficulty moving to other types of work, but the estimates do not identify the source of that difference.

5 Robustness

One possible alternative explanation for the wage results is the adjustment to remote work following the COVID-19 pandemic. Occupations differed considerably in their ability to be performed remotely, and the expansion of working from home produced lasting changes in how work was organized (Barrero et al., 2021). Wage changes associated with this transition could therefore overlap with the period in which firms began adopting generative AI.

Our baseline regressions address this possibility by interacting pre-ChatGPT work-from-home exposure with the post indicator or with year indicators. In OEWS, the work-from-home measure also enters the model used to construct the inverse-probability weights. We now examine whether the results depend on these controls, on the weighting procedure, on occupational aggregation, or on the measures of

exposure and wages.

5.1 Robustness in the OEWS Sample

Table 10 reports four variations of the pooled OEWS wage specification. Every column includes occupation-by-industry and industry-by-year fixed effects. Column (1) aggregates occupations to the three-digit SOC level while retaining four-digit NAICS industries. The coefficient remains negative and statistically significant, at -0.041 . The wage result is therefore not driven by idiosyncratic variation among narrowly defined six-digit occupations.

Columns (2) and (3) change the weighting procedure. Using weights from a parametric treatment-assignment model produces a coefficient of -0.116 . Removing the inverse-probability weights altogether yields a smaller estimate of -0.021 , which remains statistically significant. The magnitude is consequently sensitive to the way occupation-by-industry cells are weighted, but the negative sign does not depend on the baseline nonparametric weights.

Column (4) removes the work-from-home control. Rather than disappearing, the GenAI coefficient becomes more negative, moving from -0.078 in the baseline specification to -0.094 . The work-from-home adjustment therefore does not generate the negative wage relationship.

These exercises show that the OEWS result survives broader occupational aggregation and several different approaches to weighting and remote-work exposure. At the same time, the variation in magnitude across weighting schemes indicates that the composition of occupation-by-industry cells matters for the size of the estimated effect.

5.2 Robustness in the Revelio Sample

Table 11 reports the corresponding robustness exercises using Revelio. The first two columns use alternative measures of exposure and wages. Column (1) replaces predetermined cell-level exposure with observed AI exposure from the Anthropic Economic Index (Appel et al., 2025). The coefficient is -0.041 and statistically significant at the five-percent level. Because observed AI use may respond to wages, worker composition, and other labor market outcomes, we view this estimate as descriptive evidence rather than a causal effect of adoption.

Column (2) replaces Revelio’s modeled starting wages with wages advertised directly in job post-

ings. The estimate remains negative and statistically significant, at -0.105 . Appendix Figure A5 shows the corresponding dynamic estimates. This result is important because posted wages are observed independently of Revelio’s starting-wage model. The wage decline is therefore present in firms’ advertised offers as well as in our broader baseline wage measure.

Columns (3) and (4) address weighting and wage coverage. Giving every firm-by-occupation-by-MSA cell equal weight reduces the coefficient to -0.060 , but it remains negative and statistically significant. Requiring starting wages to be observed in every year produces a larger estimate of -0.188 . Neither the baseline position weights nor changes in the coverage of starting wages account for the result.

Column (5) removes the work-from-home control. The coefficient is -0.130 , compared with -0.140 in the baseline specification. As in OEWS, controlling for occupations’ pre-existing suitability for remote work has little effect on the estimate.

The final three columns address a more difficult concern: highly exposed occupations may have experienced other unusual shocks after 2022. Column (6) adds six-digit occupation-by-year fixed effects, which absorb every national shock affecting a detailed occupation in a given year. The coefficient remains negative and significant, at -0.220 . Identification now comes from differences in detailed O*NET composition across firm-by-occupation-by-MSA cells within the same six-digit occupation and year.

Column (7) includes occupation-by-MSA-by-year fixed effects. These absorb local changes in demand for a detailed occupation, including occupation-specific post-pandemic adjustments within an MSA. The coefficient is -0.211 and remains statistically significant.

Column (8) combines firm-by-occupation-by-year and occupation-by-MSA-by-year fixed effects. This specification absorbs shocks to an occupation within a particular firm and shocks to the same occupation within a particular local labor market. Even under this substantially narrower comparison, the coefficient remains negative and significant, at -0.262 .

In columns (6)-(8), the work-from-home interaction is absorbed by the occupation-by-year fixed effects rather than estimated separately. These specifications substantially reduce the concern that the baseline estimate reflects national, local, or firm-specific occupation shocks unrelated to GenAI. They cannot, however, rule out shocks occurring at a still finer level and correlated with the predetermined exposure measure.

Appendix Table [A2](#) provides complementary evidence using the GenAI-related skills reported in workers' profiles. We first divide firm-by-occupation-by-MSA cells according to whether their GenAI worker share is ever positive during the sample period. The post-ChatGPT exposure coefficient is -0.065 among cells that never report a positive share and -0.172 among cells in which GenAI-related skills are observed. The negative relationship is therefore present in both groups but is considerably stronger in the second. This comparison is descriptive because cells are classified using skill information from the complete 2018-2026 period.

The remaining specifications distinguish the wage associated with an individual worker's skills from the relationship between wages and the prevalence of those skills in the cell. Workers who report GenAI-related skills receive a conditional starting-wage premium of approximately 8-9 percent. At the same time, the coefficient on the cell-level GenAI worker share is -0.151 when entered without predetermined exposure and -0.122 when both measures are included. The post-ChatGPT exposure coefficient also remains negative after controlling for workers' skill status.

There is no contradiction between these findings. A worker who reports GenAI-related skills may command a premium relative to otherwise comparable workers, while starting wages decline in cells where those skills become more prevalent. The former is a worker-level wage differential; the latter reflects changes in the wage structure of the labor market cell.

These skill measures come from professional profiles and may respond to wages, worker selection, occupational mobility, or changes in task assignment. We therefore do not interpret the estimates as causal returns to acquiring GenAI skills or as causal effects of skill diffusion. Their more limited implication is that the baseline wage relationship cannot be reduced to the individual wage premium associated with reported GenAI skills.

Across the robustness exercises, the negative wage relationship appears in both OEWS and Revelio under alternative weighting procedures, occupational aggregations, wage measures, exposure measures, samples, and fixed-effect structures. Its magnitude changes across specifications, particularly when the weighting scheme or identifying variation changes, but its sign is remarkably stable. The result also remains when advertised wages replace modeled starting wages and when the identifying comparison is restricted to cells within the same occupation, firm, local labor market, and year.

6 Conclusion

The early labor market response to generative AI is visible clearly in wages and worker mobility. Across two different datasets, more exposed occupations experience larger wage declines after the release of ChatGPT. Thus, the absence of a broad employment decline does not mean that exposed labor markets have remained unchanged. Instead, we find large declines in the separation rate, and especially in job-to-job separations. These results describe a labor market in which outside opportunities deteriorate before existing jobs disappear. Fewer opportunities to move between employers can reduce wage pressure even when employment relationships remain intact. In this sense, weaker labor demand and greater scope for firms to set wages are not competing explanations. If demand for exposed tasks falls, alternative offers may become scarcer, mobility may slow, and workers may become less willing to leave in response to wages.

The separation-elasticity estimates are consistent with this interpretation. After ChatGPT, separations become less responsive to posted wages, particularly toward the top of the GenAI exposure distribution. The same high-exposure pattern appears among junior, mid-level, and senior workers. Applying the inverse-elasticity rule, this decline in elasticities maps into larger wage markdowns.

Within firms, junior positions experience the largest and most persistent declines, followed by mid-level positions, while the senior estimates are smaller and less precise. This gradient cannot be reduced to worker age. Wage declines appear among junior workers in every age group, including those older than 40. The decline in separation rates is broader still, extending across age and seniority groups even where the wage response is small.

Our design identifies differential changes between more and less exposed units. It does not measure changes common to all occupations, nor does the exposure index tell us which firms have adopted GenAI or how intensively they use it. The estimates also cover only the first few years following the release of ChatGPT, a period in which firms and workers are still learning how to use the technology and reorganize work. The longer-run response may therefore look different.

Still, the evidence already shows that employment can remain stable while wages fall, job-to-job mobility slows, and workers' outside opportunities deteriorate. Whether this pattern persists or eventually gives way to broader employment losses remains an open question.

References

- Abowd, John M, Francis Kramarz, and David N Margolis (1999) “High wage workers and high wage firms,” *Econometrica*, 67 (2), 251–333.
- Acemoglu, Daron and Pascual Restrepo (2018) “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” Working Paper 24196, National Bureau of Economic Research, [10.3386/w24196](https://doi.org/10.3386/w24196).
- (2020) “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128 (6), 2188–2244.
- Appel, Ruth, Peter McCrory, Alex Tamkin, Miles McCain, Tyler Neylon, and Michael Stern (2025) “Anthropic Economic Index report: Uneven geographic and enterprise AI adoption,” [10.48550/arXiv.2511.15080](https://arxiv.org/abs/10.48550/arXiv.2511.15080).
- Ashenfelter, Orley, Henry S Farber, and Michael R Ransom (2010) “Modern models of monopsony in labor markets: A brief survey.”
- Autor, David, Arindrajit Dube, and Annie McGrew (2023) “The unexpected compression: Competition at work in the low wage labor market,” Technical report, National Bureau of Economic Research.
- Autor, David, Claudia Goldin, and Lawrence F. Katz (2020) “Extending the Race between Education and Technology,” *AEA Papers and Proceedings*, 110, 347–51.
- Azar, José and Ioana Marinescu (2024a) “Monopsony power in the labor market,” in *Handbook of Labor Economics*, 5, 761–827: Elsevier.
- (2024b) “Monopsony power in the labor market: From theory to policy,” *Annual Review of Economics*, 16 (1), 491–518.
- Azar, José, Ioana Marinescu, and Marshall Steinbaum (2022) “Labor market concentration,” *Journal of Human Resources*, 57 (S), S167–S199.
- Barrero, Jose Maria, Nicholas Bloom, and Steven J. Davis (2021) “Why Working from Home Will Stick,” Working Paper 28731, National Bureau of Economic Research, [10.3386/w28731](https://doi.org/10.3386/w28731).

- Bassier, Ihsaan, Arindrajit Dube, and Suresh Naidu (2022) “Monopsony in movers: The elasticity of labor supply to firm wage policies,” *Journal of Human Resources*, 57 (S), S50–s86.
- Börschlein, Erik-Benjamin, Mario Bossler, and Martin Popp (2024) “Scarce Workers, High Wages?”.
- Bossler, Mario and Martin Popp (2026) “Labor demand on a tight leash,” *ILR Review*, 79 (4), 617–651.
- Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025a) “Canaries in the Coal Mine? Six Facts about the Employment Effects of Artificial Intelligence,” *Stanford Digital Economy Working Paper*.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025b) “Generative AI at Work,” *Quarterly Journal of Economics*, 140 (2), 889–942.
- Burdett, Kenneth and Dale T Mortensen (1998) “Wage differentials, employer size, and unemployment,” *International economic review*, 257–273.
- Caldwell, Sydnee, Arindrajit Dube, and Suresh Naidu (2023) “Monopsony makes it big,” *Journal of Economic Literature*.
- Card, David, Lawrence F Katz, and Alan B Krueger (1993) “An evaluation of recent evidence on the employment effects of minimum and subminimum wages.”
- Card, David and Alan B. Krueger (1994) “Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania,” *American Economic Review*, 84 (4), 772–793.
- Depew, Briggs and Todd A. Sørensen (2013) “The elasticity of labor supply to the firm over the business cycle,” *Labour Economics*, 24, 196–204, <https://doi.org/10.1016/j.labeco.2013.08.005>.
- Dorn, David, Florian Schoner, Moritz Seebacher, Lisa Simon, and Ludger Woessmann (2025) “Multidimensional Skills on LinkedIn Profiles: Measuring Human Capital and the Gender Skill Gap,” *arXiv preprint arXiv:2409.18638*, <https://arxiv.org/abs/2409.18638>, Version 2, revised 7 May 2025.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock (2024) “GPTs are GPTs: Labor market impact potential of LLMs,” *Science*, 384 (6702), 1306–1308.
- Goldin, Claudia and Lawrence F. Katz (2008) *The Race between Education and Technology*: Harvard University Press.

- Hampole, Michael, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller (2025) “Artificial Intelligence and the Labor Market,” *NBER Working Paper No. 33509*.
- Hansen, Stephen, Peter John Lambert, Nicholas Bloom, Steven J. Davis, Raffaella Sadun, and Bledi Taska (2023) “Remote Work across Jobs, Companies, and Space,” Working Paper 31007, National Bureau of Economic Research, [10.3386/w31007](https://doi.org/10.3386/w31007).
- Heise, Sebastian, Jeremy Pearce, and Jacob P Weber (2026) “Wage growth and labor market tightness,” *Journal of Monetary Economics*, 103903.
- Hirsch, Boris, Elke J Jahn, Alan Manning, and Michael Oberfichtner (2022) “The wage elasticity of recruitment,” Technical report, Working Paper Series in Economics.
- Hirsch, Boris, Elke J. Jahn, and Claus Schnabel (2018) “Do Employers Have More Monopsony Power in Slack Labor Markets?” *ILR Review*, 71 (3), 676–704, <https://ideas.repec.org/a/sae/ilrrev/v71y2018i3p676-704.html>.
- Hosseini, Seyed Mahdi and Guy Lichtinger (2025) “Generative ai as seniority-biased technological change: Evidence from us resume and job posting data,” *Available at SSRN*.
- Hui, Xiang, Oren Reshef, and Luofeng Zhou (2024) “The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market,” *Organization Science*, 35 (5), 1977–1989.
- Iscenko, Zanna and Fabien Curto Millet (2026) “Looking for the Ladder: Is AI Impacting Entry-Level Jobs?” the american worker project working paper, Economic Innovation Group, <https://eig.org/wp-content/uploads/2026/01/TAWP-Iscenko-Millet.pdf>.
- Karahan, Fatih, Ryan Michaels, Benjamin Pugsley, Ayşegül Şahin, and Rachel Schuh (2017) “Do job-to-job transitions drive wage fluctuations over the business cycle?” *American Economic Review*, 107 (5), 353–357.
- Katz, Lawrence F. (2025) “Beyond the Race between Education and Technology,” Technical report, Harvard University, mimeo prepared for the Jackson Hole Economic Policy Symposium.

- Katz, Lawrence F. and Kevin M. Murphy (1992) "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *The Quarterly Journal of Economics*, 107 (1), 35–78.
- Lambert, Peter John and Yannick Schindler (2026) "The Broken Ladder: AI, Remote Work, and Early-Career Hiring," SSRN Working Paper Series 6787638, Social Science Research Network, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=6787638.
- Manning, Alan (2003) *Monopsony in Motion: Imperfect Competition in the Labor Market*: Princeton University Press.
- (2011) "Imperfect competition in the labor market," in *Handbook of labor economics*, 4, 973–1041: Elsevier.
- Noy, Shakked and Whitney Zhang (2023) "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence," *Science*, 381 (6654), 187–192.
- Postel-Vinay, Fabien and Jean-Marc Robin (2002) "Equilibrium wage dispersion with worker and employer heterogeneity," *Econometrica*, 70 (6), 2295–2350.
- Sokolova, Anna and Todd Sorensen (2021) "Monopsony in labor markets: A meta-analysis," *ILR Review*, 74 (1), 27–55.
- VandeHei, Jim and Mike Allen (2025) "Behind the Curtain: A White-Collar Bloodbath," *Axios*, <https://www.axios.com/2025/05/28/ai-jobs-white-collar-unemployment-anthropic>, May 28, 2025.
- Webb, Michael (2020) "The Impact of Artificial Intelligence on the Labor Market," *Working Paper*.

7 Figures




Figure 1. Distribution of Generative AI Exposure in OEWS

Notes: The figure plots a kernel density estimate of the generative AI exposure index across the OEWS estimation sample, with detailed occupation by four-digit industry cells weighted by 2022 employment. Exposure is the occupation-level measure of [Eloundou et al. \(2024\)](#), merged to detailed occupations; it is measured ex ante, before the post-2022 period. The vertical lines denote the employment-weighted mean (0.448) and median (0.520) of the index.




Figure 2. Effect of Generative AI Exposure on Wages in OEWS

Notes: The figure plots the event-study coefficients β_τ from the continuous-treatment specification in equation (3.1), estimated on the OEWS sample of detailed occupation by four-digit industry by year cells. Each coefficient is the interaction of the ex-ante generative AI exposure index (Eloundou et al., 2024) with a year indicator, measured relative to the omitted reference year 2022; the outcome is the real log median hourly wage. The specification includes occupation-by-industry and industry-by-year (four-digit NAICS by year) fixed effects, and observations are weighted by the product of stabilized nonparametric inverse-probability weights and 2022 cell employment. The specification also includes year-specific interactions with pre-ChatGPT occupational work-from-home exposure. Standard errors are clustered by three-digit industry, and the shaded region denotes the 95% confidence interval. The corresponding estimates are reported in Table 3.




Figure 3. Effect of Generative AI Exposure on Employment in OEWS

Notes: The figure plots the event-study coefficients β_τ from equation (3.1), estimated on the OEWS sample of detailed occupation by four-digit industry by year cells; the outcome is log employment. Coefficients are interactions of the ex-ante generative AI exposure index (Eloundou et al., 2024) with year indicators, relative to the omitted year 2022. The specification includes occupation-by-industry and industry-by-year (four-digit NAICS by year) fixed effects, and observations are weighted by stabilized nonparametric inverse-probability weights times 2022 cell employment. Standard errors are clustered by three-digit industry, and the shaded region denotes the 95% confidence interval. Estimates are reported in Table 3.




Figure 4. Effect of Generative AI Exposure on Wages in Revelio

Notes: The figure plots the dynamic difference-in-differences coefficients from equation (3.3), using the real log starting wage of newly created positions as the dependent variable. The coefficients interact the predetermined cell-level GenAI exposure measure, AI_{fom} , defined in Section 2.2, with year indicators and are measured relative to the omitted year 2022. The specification also includes year-specific interactions with pre-ChatGPT occupational work-from-home exposure. The regression is weighted by the average number of newly created positions in the firm-by-occupation-by-MSA cell during 2018-2022 and includes firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. The specification also includes year-specific interactions with pre-ChatGPT occupational work-from-home exposure. Standard errors are two-way clustered by firm and occupation, and the shaded region denotes a 95% confidence interval. The corresponding estimates are reported in Table 4.






Figure 5. Effect of Generative AI Exposure on the Separation Rate in Revelio

Notes: The figure plots the dynamic difference-in-differences coefficients from equation (3.3), using the separation rate as the dependent variable. The separation rate is defined as annual separations divided by lagged employment in the firm-by-occupation-by-MSA cell. The coefficients interact the predetermined cell-level GenAI exposure measure, AI_{fom} , constructed from the cell's detailed O*NET occupational composition using [Eloundou et al. \(2024\)](#), with year indicators and are measured relative to the omitted year 2022. The specification also includes year-specific interactions with occupational work-from-home exposure. The worker-flow sample is restricted to observations for which the separation rate, job-to-job transition rate, and persistent-non-employment separation rate are all no greater than one. The regression is weighted by average cell employment during 2018-2022 and includes firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Standard errors are two-way clustered by firm and occupation, and the shaded region denotes a 95% confidence interval. Estimates are reported in Table 4.



(a) Job-to-Job Separations



(b) Separations to Persistent Non-employment

Figure 6. Effects of Generative AI Exposure on Job-to-Job Transitions and Separations to Persistent Non-Employment

Notes: The figure plots the dynamic difference-in-differences coefficients from equation (3.3). Panel A uses the job-to-job transition rate as the dependent variable, while Panel B uses the persistent-non-employment separation rate. Each rate is calculated as the corresponding annual flow divided by lagged employment in the firm-by-occupation-by-MSA cell. The coefficients interact the predetermined cell-level GenAI exposure measure, AI_{fom} , constructed from the cell's detailed O*NET occupational composition using [Eloundou et al. \(2024\)](#), with year indicators and are measured relative to the omitted year 2022. The specifications also include year-specific interactions with occupational work-from-home exposure. The worker-flow sample is restricted to observations for which the separation rate, job-to-job transition rate, and persistent-non-employment separation rate are all no greater than one. Regressions are weighted by average cell employment during 2018-2022 and include firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Standard errors are two-way clustered by firm and occupation, and the shaded regions denote 95% confidence intervals. Estimates are reported in Table 4.




Figure 7. Separation Elasticities over the Generative AI Exposure Distribution

Notes: The figure reports separation elasticities before and after ChatGPT from the heterogeneous spline AKM-IV specification. The continuously estimated spline is evaluated at the mean exposure within each decile; deciles are display points rather than discrete treatment groups. Log posted wages and their interactions are instrumented using the corresponding interactions with the cross-fitted AKM firm wage component. The specification includes worker, MSA-by-occupation, MSA-by-year, and occupation-by-year fixed effects. Standard errors are clustered by MSA-by-occupation labor market. Vertical bars denote pointwise 95% confidence intervals, treating fitted separation probabilities as fixed.






Figure 8. Post-ChatGPT Change in the Absolute Separation Elasticity across the Generative AI Exposure Distribution


Notes: The figure reports the estimated change in the separation elasticity between the post- and pre-ChatGPT periods, $\Delta|\eta^S|(x) = |\eta_{\text{post}}^S(x)| - |\eta_{\text{pre}}^S(x)|$, over the distribution of occupational exposure to generative AI, x . The spline is evaluated at the mean exposure within each decile of the estimation sample. The pre-ChatGPT period includes years through 2022, while the post-ChatGPT period begins in 2023. The elasticities are obtained from the heterogeneous spline AKM-IV specification described in Section 3.3. Log posted wages and their interactions with the restricted-cubic-spline basis and the post-ChatGPT indicator are instrumented using the corresponding interactions with the cross-fitted AKM firm wage component. The specification includes worker, MSA-by-occupation labor market, MSA-by-year, and occupation-by-year fixed effects. Standard errors are clustered by MSA-by-occupation labor market, and the vertical bars denote pointwise 95% confidence intervals for the estimated post-pre difference. These confidence intervals account for uncertainty in the IV wage-slope estimates while treating the fitted separation probabilities as fixed. The horizontal dashed line indicates no change in the absolute separation elasticity. Negative values indicate a decline in workers' responsiveness to posted wages, whereas positive values indicate an increase in its absolute magnitude.



(a) Junior



(b) Mid-level



(c) Senior

Figure 9. Separation Elasticities across the Generative AI Exposure Distribution by Seniority

Notes: Each panel reports separation elasticities before and after ChatGPT for the indicated 2022 seniority group. The continuously estimated spline is evaluated at the mean exposure within each decile of the corresponding seniority sample. Log posted wages are instrumented using the cross-fitted AKM firm wage component. Vertical bars denote pointwise 95% confidence intervals. The specifications include worker, MSA-by-occupation, MSA-by-year, and occupation-by-year fixed effects. Standard errors are clustered by MSA-by-occupation labor market, and fitted separation probabilities are treated as fixed when constructing the confidence intervals.




Figure 10. Effect of Generative AI Exposure on Wages by Seniority

Notes: Each panel plots dynamic difference-in-differences coefficients from equation (3.3), estimated separately by seniority. The dependent variable is the real log starting wage. Coefficients interact the predetermined cell-level GenAI exposure measure, AI_{fom} , defined in Section 2.2, with year indicators, relative to 2022. All specifications control for year-specific interactions with pre-ChatGPT work-from-home exposure. Regressions include firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects, weighted by the average number of newly created positions in the corresponding seniority group during 2018–2022. Standard errors are two-way clustered by firm and occupation; shaded regions denote 95% confidence intervals (Table 7).






Figure 11. Distribution of Age by Seniority Level

Notes: The figure plots kernel density estimates of worker age within each seniority group (junior, mid-level, or senior) in the Revelio analysis sample. Seniority is measured using Revelio's ordinal hierarchy scale, aggregated into three groups as described in Section 2. The figure illustrates the overlap in the age distributions across seniority groups.



(a) Wages



(b) Separations

Figure 12. Effects of Generative AI Exposure on Wages and Separations by Seniority and Age

Notes: The figure reports pooled difference-in-differences coefficients on the interaction between the post-ChatGPT indicator and predetermined cell-level GenAI exposure, AI_{fom} , estimated separately by seniority group (junior, mid-level, and senior) and age group (22-30, 31-40, and 41-70). Panel A uses the real log starting wage of newly created positions as the dependent variable. Because the dependent variable is in logarithms, the wage coefficients are semi-elasticities. Panel B uses the seniority-by-age separation rate, defined as the number of annual separations in the corresponding group divided by its lagged employment stock. The separation coefficients are reported in levels and measure the change in the separation rate associated with a one-unit increase in GenAI exposure; they are not scaled by the mean of the dependent variable. Observations with missing or nonpositive lagged group employment and observations for which the corresponding seniority-by-age separation rate exceeds one are excluded from Panel B. Wage regressions are weighted by the average number of newly created positions in the corresponding seniority-by-age group during 2018-2022, while separation regressions are weighted by the corresponding groups average employment during 2018-2022. All specifications jointly control for the interaction between pre-ChatGPT occupational work-from-home exposure and the post-ChatGPT indicator and include firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Standard errors are two-way clustered by firm and occupation, and the vertical bars denote 95% confidence intervals.

8 Tables

Table 1. Summary Statistics: OEWS Analysis Sample

Variable	Mean	N	SD	Max	Min
Log real hourly wage	2.308	24,736	0.518	3.676	0.897
Real hourly wage (USD)	11.566	24,736	6.516	39.499	2.453
Log employment	10.559	24,736	1.819	13.691	3.401
Employment	140,285.844	24,736	221,112.390	883,360.000	30.000
AI exposure	0.448	24,736	0.226	1.000	0.000
WFH exposure	0.022	24,736	0.018	0.179	0.000

Notes: The table reports weighted summary statistics for detailed six-digit SOC occupation-by-four-digit NAICS industry-by-year observations in knowledge-intensive industries, defined as NAICS two-digit sectors 51-56. The sample is a balanced panel covering 2018-2025. Statistics are weighted by occupation-by-industry employment measured in 2022, so that post-ChatGPT employment changes do not alter the relative importance of cells in the sample. Wage measures report median hourly wages from OEWS, deflated using the Consumer Price Index and expressed in real terms. Log real hourly wages and log employment are constructed by taking the natural logarithm of the corresponding levels. AI exposure is the occupation-level measure of ex-ante generative AI exposure from [Eloundou et al. \(2024\)](#). WFH exposure is the pre-ChatGPT occupation-level prevalence of remote-work job advertisements constructed from Revelio postings following [Hansen et al. \(2023\)](#).

Table 2. Summary Statistics: Revelio Analysis Sample

Variable	Mean	N	SD	Max	Min
Real log wage	10.340	5,746,883	0.503	12.801	3.245
Employment	210.051	15,218,901	1,354.145	26,219.000	1.000
Separation rate	0.147	15,042,765	0.169	1.000	0.000
Job-to-job transition rate	0.084	15,042,765	0.123	1.000	0.000
Persistent-non-employment separation rate	0.057	15,042,765	0.095	1.000	0.000
New positions	47.380	15,218,901	326.842	7,488.000	0.000
Ending positions	41.115	15,218,901	273.936	6,880.000	0.000
AI exposure	0.565	15,218,901	0.153	1.000	0.000
WFH exposure	0.048	15,218,901	0.024	0.334	0.000
GenAI worker share	0.005	15,218,901	0.026	1.000	0.000

Notes: The table reports weighted summary statistics for firm-occupation-MSA-year observations in knowledge-intensive industries, defined as NAICS two-digit sectors 51-56. The sample is a balanced panel covering 2018-2026 and is restricted to cells with more than one employee in 2022. Wages are real log wages for newly created positions and are weighted by the average number of new positions in the cell during 2018-2022. Employment, worker-flow rates, shares, and exposure measures are weighted by average employment during 2018-2022. The separation rate is defined as annual separations divided by lagged cell employment. The job-to-job transition rate and persistent-non-employment separation rate are defined analogously, using job-to-job transitions and separations without observed re-employment within 90 days as their respective numerators. Observations with missing or nonpositive lagged employment are excluded from the corresponding rate statistic. Observations for which any of the three worker-flow rates exceeds one are excluded from all three rate statistics, ensuring a common sample across the worker-flow outcomes. This restriction does not affect the descriptive statistics reported for the remaining variables. New and ending positions are counts measured at the firm-occupation-MSA-year level. The GenAI worker share is the share of active workers associated with GenAI-related skills reported in professional profiles. It is a proxy for the prevalence of GenAI-related skills and human capital within the cell, not a direct measure of generative AI use or adoption. AI exposure is the predetermined cell-level measure constructed from the detailed O*NET occupational composition of the firm-occupation-MSA cell on November 30, 2022. WFH exposure is the pre-ChatGPT occupation-level prevalence of remote-work job advertisements.

Table 3. Effects of Generative AI and Work-from-Home Exposure on Wages and Employment in OEWS

	Log Wage		Log Employment	
	DiD (1)	Dynamic DiD (2)	DiD (3)	Dynamic DiD (4)
<i>Panel A: Generative AI exposure</i>				
Post-ChatGPT $\times$ AI exposure	-0.078*** (0.016)		-0.188 (0.177)	
2018 $\times$ AI exposure		0.016 (0.018)		0.142 (0.177)
2019 $\times$ AI exposure		0.016 (0.023)		0.213 (0.193)
2020 $\times$ AI exposure		-0.004 (0.028)		0.197 (0.208)
2021 $\times$ AI exposure		-0.005 (0.020)		0.230* (0.123)
2023 $\times$ AI exposure		-0.036 (0.028)		-0.117 (0.091)
2024 $\times$ AI exposure		-0.114*** (0.031)		-0.365 (0.287)
2025 $\times$ AI exposure		-0.083** (0.037)		-0.081 (0.165)
<i>Panel B: Work-from-home exposure</i>				
Post-ChatGPT $\times$ WFH exposure	-0.437* (0.236)		1.202*** (0.362)	
2018 $\times$ WFH exposure		0.508 (0.466)		-1.693 (1.611)
2019 $\times$ WFH exposure		0.265 (0.492)		-0.597 (1.431)
2020 $\times$ WFH exposure		-0.061 (0.515)		0.415 (1.041)
2021 $\times$ WFH exposure		-0.272 (0.769)		-0.114 (0.403)
2023 $\times$ WFH exposure		0.071 (0.239)		0.549 (0.381)
2024 $\times$ WFH exposure		-0.521 (0.364)		2.353** (1.010)
2025 $\times$ WFH exposure		-0.860*** (0.157)		0.705 (0.867)
Observations	12,368	24,736	12,368	24,736
R-squared	0.987	0.982	0.994	0.988
Occupation-by-industry FE	Yes	Yes	Yes	Yes
Industry-by-year FE	Yes	Yes	Yes	Yes
Sample	2022-2025	2018-2025	2022-2025	2018-2025

Notes: The table reports continuous-treatment difference-in-differences estimates using the OEWS occupation-by-industry panel. The dependent variable is the real log median hourly wage in columns (1) and (2) and log employment in columns (3) and (4). Odd-numbered columns report pooled post-ChatGPT effects estimated over 2022-2025. Even-numbered columns report year-specific effects using the full 2018-2025 panel, with 2022 serving as the omitted reference year. The pooled specifications interact generative AI exposure and pre-ChatGPT occupational work-from-home exposure with an indicator equal to one from 2023 onward. The dynamic specifications interact both exposure measures with year indicators. All regressions include occupation-by-industry and four-digit industry-by-year fixed effects. Observations are weighted by the product of the stabilized nonparametric inverse-probability weight and occupation-by-industry employment measured in 2022. Standard errors are clustered by three-digit NAICS industry and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4. Generative AI Exposure, Work-from-Home Exposure, Wages, and Worker Flows

	Log Wage		Separation Rate		Job-to-Job Rate		Persist. Non-Employment Rate	
	DiD (1)	Dynamic DiD (2)	DiD (3)	Dynamic DiD (4)	DiD (5)	Dynamic DiD (6)	DiD (7)	Dynamic DiD (8)
<i>Panel A: Generative AI exposure</i>								
Post-ChatGPT $\times$ AI exposure	-0.140*** (0.032)		-0.074*** (0.014)		-0.069*** (0.012)		-0.005 (0.006)	
2018 $\times$ AI exposure		-0.013 (0.021)		0.006 (0.010)		-0.003 (0.009)		0.007 (0.006)
2019 $\times$ AI exposure		-0.002 (0.019)		0.012 (0.010)		0.017** (0.008)		-0.006 (0.005)
2020 $\times$ AI exposure		-0.007 (0.018)		0.003 (0.011)		0.008 (0.009)		-0.005 (0.005)
2021 $\times$ AI exposure		-0.008 (0.007)		-0.004 (0.007)		-0.002 (0.005)		-0.003 (0.003)
2023 $\times$ AI exposure		-0.092*** (0.024)		-0.048*** (0.010)		-0.046*** (0.010)		-0.002 (0.003)
2024 $\times$ AI exposure		-0.098*** (0.023)		-0.071*** (0.014)		-0.067*** (0.012)		-0.005 (0.005)
2025 $\times$ AI exposure		-0.187*** (0.040)		-0.076*** (0.015)		-0.071*** (0.012)		-0.007 (0.007)
2026 $\times$ AI exposure		-0.266*** (0.055)		-0.088*** (0.018)		-0.076*** (0.015)		-0.013** (0.007)
<i>Panel B: Work-from-home exposure</i>								
Post-ChatGPT $\times$ WFH exposure	0.422** (0.177)		0.039 (0.069)		0.047 (0.034)		-0.014 (0.039)	
2018 $\times$ WFH exposure		0.151 (0.167)		0.108* (0.057)		0.059 (0.059)		0.047 (0.037)
2019 $\times$ WFH exposure		0.078 (0.146)		0.081* (0.046)		0.041 (0.040)		0.043** (0.022)
2020 $\times$ WFH exposure		0.059 (0.187)		0.095** (0.045)		0.036 (0.055)		0.060** (0.025)
2021 $\times$ WFH exposure		0.023 (0.061)		0.054 (0.034)		0.037 (0.033)		0.022 (0.022)
2023 $\times$ WFH exposure		0.384*** (0.115)		0.082 (0.055)		0.055 (0.041)		0.021 (0.020)
2024 $\times$ WFH exposure		0.460*** (0.111)		0.091 (0.068)		0.079* (0.048)		0.006 (0.034)
2025 $\times$ WFH exposure		0.531*** (0.189)		0.062 (0.083)		0.070 (0.056)		-0.009 (0.045)
2026 $\times$ WFH exposure		0.603** (0.275)		0.191* (0.112)		0.123 (0.076)		0.065 (0.048)
Observations	4,896,686	4,896,686	13,506,910	13,506,910	13,506,910	13,506,910	13,506,910	13,506,910
R-squared	0.891	0.891	0.483	0.483	0.443	0.443	0.355	0.356
Firm-by-occupation-by-MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Sample	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026

Notes: The table reports continuous-treatment difference-in-differences estimates using the Revelio firm-by-occupation-by-MSA panel. The wage outcome is the real log starting wage of newly created positions. The worker-flow outcomes are the separation rate, the job-to-job transition rate, and the persistent-non-employment separation rate, each calculated as the corresponding annual flow divided by lagged cell employment. The worker-flow sample excludes observations for which any of the three rates exceeds one, ensuring a common estimation sample across the worker-flow outcomes. Odd-numbered columns report pooled post-ChatGPT effects, while even-numbered columns report year-specific effects relative to 2022. The pooled specifications interact predetermined cell-level GenAI exposure and occupational work-from-home exposure with a post-ChatGPT indicator. The dynamic specifications interact the cell-level GenAI exposure measure and the occupation-level work-from-home exposure measure with year indicators, with 2022 serving as the omitted reference year. Wage regressions are weighted by the average number of new positions during 2018-2022, while worker-flow regressions are weighted by average employment during 2018-2022. All specifications include firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. The generative AI treatment is the predetermined cell-level exposure measure, $AI_{i,t}^{dom}$, constructed from the detailed O*NET occupational composition of the firm-occupation-MSA cell on November 30, 2022. Standard errors are two-way clustered by firm and occupation and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5. Generative AI Exposure and Wages: Sequential Demographic Controls

	Baseline (1)	Demographics (2)	College (3)	Experience (4)
Post-ChatGPT $\times$ AI exposure	-0.140*** (0.032)	-0.140*** (0.031)	-0.142*** (0.031)	-0.132*** (0.027)
Post-ChatGPT $\times$ WFH exposure	0.422** (0.177)	0.425** (0.176)	0.428** (0.176)	0.317** (0.156)
Female share		-0.045*** (0.004)	-0.044*** (0.004)	-0.039*** (0.003)
Black share		-0.040*** (0.002)	-0.039*** (0.002)	-0.020*** (0.002)
Hispanic share		-0.057*** (0.003)	-0.055*** (0.003)	-0.028*** (0.002)
Asian/Pacific Islander share		-0.033*** (0.004)	-0.034*** (0.004)	0.006 (0.004)
College share			0.049*** (0.007)	0.052*** (0.006)
Average experience				0.014*** (0.001)
Observations	4,896,686	4,893,682	4,893,682	4,893,682
R-squared	0.891	0.892	0.892	0.901
Dependent variable mean	10.26	10.26	10.26	10.26
Demographic sample	No	Yes	Yes	Yes
Firm-by-occupation-by-MSA FE	Yes	Yes	Yes	Yes
Firm-by-year FE	Yes	Yes	Yes	Yes
MSA-by-year FE	Yes	Yes	Yes	Yes
Sample	2018-2026	2018-2026	2018-2026	2018-2026

Notes: The table reports continuous-treatment difference-in-differences estimates of the relationship between predetermined cell-level GenAI exposure and real log starting wages in Revelio. Column (1) presents the baseline specification using the full analysis sample. Column (2) restricts the analysis to observations matched to the demographic data and controls for female, Black, Hispanic, and Asian/Pacific Islander employment shares. Column (3) additionally controls for the college-educated share. Column (4) further controls for average worker experience. All specifications control for the interaction between occupational work-from-home exposure and the post-ChatGPT indicator and include firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Regressions are weighted by average new positions during 2018-2022. Standard errors are two-way clustered by firm and occupation and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6. Heterogeneous Effects of Generative AI Exposure on Wages

<i>Panel A: By industry</i>						
	Information (1)	Finance (2)	Real estate (3)	Professional services (4)	Management (5)	Administrative services (6)
Post-ChatGPT $\times$ AI exposure	−0.195*** (0.039)	−0.127*** (0.032)	−0.014 (0.040)	−0.092*** (0.031)	−0.030 (0.072)	0.030 (0.043)
Post-ChatGPT $\times$ WFH exposure	0.335 (0.278)	0.474** (0.193)	0.267 (0.310)	0.215 (0.208)	0.104 (0.580)	0.111 (0.373)
Observations	1,224,956	1,410,891	309,154	1,552,124	9,050	389,406
R-squared	0.902	0.879	0.850	0.869	0.840	0.903
Dependent variable mean	10.37	10.24	10.12	10.29	10.00	10.01
Firm-by-occupation-by-MSA FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
MSA-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026

<i>Panel B: By firm size in 2022</i>			
	Fewer than 50 employees (1)	50-250 employees (2)	More than 250 employees (3)
Post-ChatGPT $\times$ AI exposure	−0.066** (0.028)	−0.082*** (0.023)	−0.146*** (0.033)
Post-ChatGPT $\times$ WFH exposure	0.378*** (0.119)	0.228* (0.119)	0.441** (0.189)
Observations	450,095	1,073,565	3,372,704
R-squared	0.906	0.877	0.888
Dependent variable mean	10.11	10.20	10.30
Firm-by-occupation-by-MSA FE	Yes	Yes	Yes
Firm-by-year FE	Yes	Yes	Yes
MSA-by-year FE	Yes	Yes	Yes
Sample	2018-2026	2018-2026	2018-2026

Notes: The table reports heterogeneous continuous-treatment difference-in-differences estimates of the relationship between predetermined cell-level GenAI exposure and real log starting wages in Revelio. Panel A estimates the specification separately by two-digit NAICS industry. Panel B divides firms according to their total employment in 2022, aggregated across occupation-by-MSA cells. The reported coefficients correspond to the interactions between generative AI exposure and occupational work-from-home exposure, respectively, and an indicator equal to one from 2023 onward. All specifications include firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Regressions are weighted by the average number of newly created positions in each cell during 2018-2022. The sample covers 2018-2026, is restricted to a balanced panel of cells in knowledge-intensive industries, and requires each cell to have more than one employee in 2022. Standard errors are two-way clustered by firm and occupation and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7. Generative AI Exposure and Wages by Seniority in Revelio

	Junior		Mid-level		Senior	
	DiD (1)	Dynamic DiD (2)	DiD (3)	Dynamic DiD (4)	DiD (5)	Dynamic DiD (6)
<i>Panel A: Generative AI exposure</i>						
Post-ChatGPT × AI exposure	−0.100*** (0.029)		−0.060*** (0.017)		−0.043 (0.031)	
2018 × AI exposure		0.005 (0.020)		0.006 (0.023)		−0.010 (0.024)
2019 × AI exposure		0.009 (0.017)		0.025 (0.015)		−0.013 (0.022)
2020 × AI exposure		0.003 (0.018)		0.020 (0.015)		0.000 (0.019)
2021 × AI exposure		−0.006 (0.006)		−0.004 (0.007)		0.002 (0.010)
2023 × AI exposure		−0.058*** (0.019)		−0.060*** (0.011)		−0.069*** (0.017)
2024 × AI exposure		−0.080*** (0.028)		−0.040*** (0.012)		−0.040* (0.022)
2025 × AI exposure		−0.129*** (0.041)		−0.036** (0.017)		−0.031 (0.037)
2026 × AI exposure		−0.181*** (0.057)		−0.075*** (0.024)		−0.043 (0.039)
<i>Panel B: Work-from-home exposure</i>						
Post-ChatGPT × WFH exposure	0.201 (0.139)		0.373** (0.165)		−0.088 (0.163)	
2018 × WFH exposure		0.119 (0.183)		0.147 (0.191)		0.085 (0.163)
2019 × WFH exposure		0.046 (0.173)		0.040 (0.180)		0.054 (0.115)
2020 × WFH exposure		0.059 (0.159)		0.201 (0.207)		0.057 (0.146)
2021 × WFH exposure		0.119** (0.047)		−0.020 (0.094)		−0.070 (0.066)
2023 × WFH exposure		0.139* (0.079)		0.181*** (0.068)		0.136 (0.100)
2024 × WFH exposure		0.312** (0.150)		0.290 (0.191)		−0.012 (0.175)
2025 × WFH exposure		0.419* (0.217)		0.561* (0.324)		−0.210 (0.225)
2026 × WFH exposure		0.188 (0.223)		1.062** (0.505)		−0.281 (0.326)
Observations	3,022,211	3,022,211	1,866,737	1,866,737	1,154,855	1,154,855
R-squared	0.924	0.924	0.847	0.847	0.704	0.704
Firm-by-occupation-by-MSA FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
MSA-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
Sample	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026

Notes: The table reports continuous-treatment difference-in-differences estimates of the relationship between predetermined cell-level GenAI exposure and real log starting wages in Revelio, separately for junior, mid-level, and senior positions. Odd-numbered columns report pooled interactions between each exposure measure and an indicator for the post-ChatGPT period. Even-numbered columns report year-specific interactions, with 2022 serving as the omitted reference year. Panel A reports the coefficients associated with generative AI exposure, while Panel B reports the coefficients associated with occupational work-from-home exposure. All specifications include firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Regressions are weighted by the average number of newly created positions for the corresponding seniority group during 2018-2022. The sample covers 2018-2026 and is restricted to a balanced panel of firm-by-occupation-by-MSA cells in knowledge-intensive industries with more than one employee in 2022. Standard errors are two-way clustered by firm and occupation and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8. Linear Posted-Wage Separation Slopes and Elasticities

	Overall (1)	Junior (2)	Mid-level (3)	Senior (4)
Log posted wage	−0.1880*** (0.0125)	−0.2088*** (0.0210)	−0.3198*** (0.0176)	−0.2944*** (0.0270)
Log posted wage × centered AI exposure	−0.4152*** (0.0679)	−0.5413*** (0.1097)	−0.6440*** (0.0924)	−0.6043*** (0.1130)
Log posted wage × Post-ChatGPT	0.0962*** (0.0086)	0.0933*** (0.0139)	0.1841*** (0.0129)	0.1850*** (0.0189)
Log posted wage × centered AI exposure × Post-ChatGPT	0.2602*** (0.0444)	0.2803*** (0.0651)	0.5332*** (0.0857)	0.4704*** (0.0901)
Average separation elasticity	−2.272 (0.145)	−2.528 (0.231)	−3.661 (0.193)	−3.228 (0.315)
Kleibergen-Paap rk <i>F</i> -statistic	194.03	184.90	229.25	67.64
Observations	34,371,006	11,221,930	7,319,798	5,282,449
Worker FE	Yes	Yes	Yes	Yes
MSA-by-occupation FE	Yes	Yes	Yes	Yes
MSA-by-year FE	Yes	Yes	Yes	Yes
Occupation-by-year FE	Yes	Yes	Yes	Yes

Notes: The table reports AKM-IV estimates from the linear posted-wage separation specification. The dependent variable is an indicator for whether a worker separates from the firm. Column (1) uses the full worker sample, while columns (2)-(4) estimate the model separately according to workers' seniority in 2022. Workers without a valid 2022 seniority classification and workers in Revelio's executive category are excluded from the seniority-specific samples. Posted wages and all interactions containing posted wages are instrumented using the corresponding interactions with the cross-fitted AKM firm wage component. The Post-ChatGPT indicator equals one beginning in 2023. AI exposure is centered at its mean in the relevant estimation sample; the seniority-specific specifications use a common mean across the three seniority groups. The average separation elasticity divides the wage slope averaged over the full estimation period by the average fitted separation probability. The average wage slope evaluates centered AI exposure, the Post-ChatGPT indicator, and their interaction at their respective sample means. Fitted separation probabilities are obtained from separate logit models and are treated as fixed when calculating the elasticity standard errors. Thus, the reported elasticity is the ratio of the average wage slope to the average fitted separation probability, rather than the average of worker-specific elasticities. All specifications include worker, MSA-by-occupation, MSA-by-year, and occupation-by-year fixed effects. Standard errors are clustered by MSA-by-occupation labor market and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9. Implied Relative Wage Markdowns at Generative AI Exposure Quintile Means

Exposure quintile	$ \eta^S $	ε^L	ν	$\% \Delta \nu$
<i>Panel A: Pre-ChatGPT period</i>				
Q1 (0.39)	1.874	3.748	0.267	-
Q2 (0.50)	1.998	3.996	0.250	-
Q3 (0.57)	2.288	4.576	0.219	-
Q4 (0.65)	2.646	5.292	0.189	-
Q5 (0.84)	5.158	10.316	0.097	-
<i>Panel B: Post-ChatGPT period</i>				
Q1 (0.39)	1.356	2.712	0.369	38.2
Q2 (0.50)	1.318	2.635	0.379	51.6
Q3 (0.57)	1.438	2.876	0.348	59.1
Q4 (0.65)	1.618	3.237	0.309	63.5
Q5 (0.84)	2.947	5.893	0.170	75.1

Notes: The table reports spline-based separation elasticities, implied firm-level labor-supply elasticities, and implied relative wage markdowns evaluated at the mean exposure within each quintile. The spline is estimated continuously; quintiles define evaluation points rather than discrete treatment groups. The value in parentheses beside each quintile denotes the corresponding mean exposure. Panel A covers the pre-ChatGPT period through 2022, while Panel B covers the post-ChatGPT period beginning in 2023. The absolute separation elasticities, $|\eta^S|$, are obtained from the spline AKM-IV specification evaluated at the mean exposure of each quintile. Log posted wages and their interactions are instrumented using the corresponding interactions with the cross-fitted AKM firm wage component. Following [Manning \(2003\)](#), the firm-level labor-supply elasticity is approximated as $\varepsilon^L = 2|\eta^S|$. Under the canonical inverse-elasticity benchmark, the implied wage markdown relative to the wage is $\nu = \frac{MRPL - w}{w} = \frac{1}{\varepsilon^L}$. The column headed $\% \Delta \nu$ reports the pre-to-post percentage change in the implied markdown within each exposure quintile, calculated from the underlying unrounded estimates as $100(X^{\text{post}}/X^{\text{pre}} - 1)$. Fitted separation probabilities are treated as fixed. The reported markdowns are illustrative benchmark calibrations and should not be interpreted as unrestricted estimates of firms' markdowns or workers' marginal revenue products.

Table 10. Robustness of the Effect of Generative AI Exposure on Wages in OEWS

	NAICS-4 $\times$ SOC-3 (1)	Parametric IPW (2)	No IPW (3)	No WFH Control (4)
Post-ChatGPT $\times$ AI exposure	−0.041** (0.015)	−0.116* (0.059)	−0.021** (0.009)	−0.094*** (0.024)
Post-ChatGPT $\times$ WFH exposure	−0.529 (0.583)	−0.709** (0.256)	−0.276** (0.104)	
Observations	5,020	12,368	12,368	12,368
R-squared	0.990	0.990	0.995	0.988
Occupation-by-industry FE	Yes	Yes	Yes	Yes
Industry-by-year FE	Yes	Yes	Yes	Yes
Sample	2022-2025	2022-2025	2022-2025	2022-2025

Notes: The table reports wage robustness exercises using OEWS data. The dependent variable is the real log median hourly wage. All columns restrict the sample to 2022-2025 and estimate a pooled post-ChatGPT effect. All specifications include occupation-by-industry and industry-by-year fixed effects. Column (1) aggregates the analysis to the four-digit NAICS industry by three-digit SOC occupation level. Column (2) uses a parametric specification for the inverse-probability weights (IPW). Column (3) reports estimates without inverse-probability weighting. Column (4) reports the baseline nonparametric-IPW specification without controlling for postings-based work-from-home exposure. Columns (1)-(3) control for the interaction between pre-ChatGPT occupational work-from-home exposure and the post-ChatGPT indicator; the corresponding coefficients are reported in columns (1)-(3). Standard errors are clustered by three-digit NAICS industry and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11. Robustness of the Effect of Generative AI Exposure on Wages in Revelio

	(1) Observed AI exposure	(2) Posted wages	(3) Unweighted	(4) Balanced wage panel	(5) No WFH control	(6) Alternative FE (a)	(7) Alternative FE (b)	(8) Alternative FE (c)
Post-ChatGPT $\times$ AI exposure		-0.105*** (0.040)	-0.060*** (0.022)	-0.188*** (0.035)	-0.130*** (0.033)	-0.220*** (0.058)	-0.211*** (0.044)	-0.262*** (0.067)
Post-ChatGPT $\times$ Observed AI exposure	-0.041** (0.021)							
Post-ChatGPT $\times$ WFH exposure	0.250 (0.174)	1.216** (0.503)	0.267* (0.137)	0.514** (0.219)				
Observations	4,670,733	2,506,558	4,896,686	675,531	4,896,686	4,896,204	4,702,947	2,617,662
R-squared	0.893	0.762	0.817	0.927	0.891	0.893	0.898	0.937
Firm-by-occupation-by-MSA FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
MSA-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Occupation-by-year FE	No	No	No	No	No	Yes	Yes	Yes
Occupation-by-MSA-by-year FE	No	No	No	No	No	No	Yes	Yes
Firm-by-occupation-by-year FE	No	No	No	No	No	No	No	Yes
WFH control	Yes	Yes	Yes	Yes	No	Absorbed	Absorbed	Absorbed
Sample	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026	2018-2026

Notes: The table reports robustness estimates for real log wages in Revelio. Columns (1) and (3)-(8) use the real log starting wage of newly created positions, while column (2) uses real log posted wages. Column (1) replaces the baseline predetermined cell-level AI exposure measure with observed AI exposure from the Anthropic Economic Index. Column (2) uses posted wages from job advertisements and is weighted by average vacancies during 2018-2022. Column (3) gives every firm-occupation-MSA cell equal weight, column (4) requires starting wages to be observed in every year, and column (5) omits the work-from-home control. Columns (6)-(8) examine alternative fixed-effect structures. Column (6) includes six-digit occupation-by-year fixed effects; column (7) includes occupation-by-MSA-by-year fixed effects; and column (8) includes both firm-by-occupation-by-year and occupation-by-MSA-by-year fixed effects. In columns (6)-(8), the interaction between occupational work-from-home exposure and the post-ChatGPT period is absorbed by the indicated occupation-by-year or occupation-by-MSA-by-year fixed effects. All specifications include firm-by-occupation-by-MSA fixed effects. The remaining fixed effects are indicated in the table. Except for the unweighted specification in column (3), the starting-wage regressions are weighted by the average number of new positions during 2018-2022; the posted-wage regression in column (2) is weighted by average vacancies over the same period. Columns (1)-(4) include the interaction between pre-ChatGPT occupational work-from-home exposure and the post-ChatGPT indicator explicitly, column (5) omits it, and columns (6)-(8) absorb it through the indicated fixed effects. The observed-exposure specification in column (1) should be interpreted as a complementary descriptive exercise because observed AI use may respond endogenously to wages and worker composition. Standard errors are two-way clustered by firm and occupation and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix

A Additional Figures and Tables




Figure A1. Correlation between OEWS and Revelio Wages

Notes: The figure is a scatterplot of average log real wages in Revelio against average log real wages in OEWS, at the detailed occupation by four-digit industry level; each point is one occupation-by-industry cell and the solid line is a linear fit. The sample-wide correlation is 0.81. Wage measures are defined in [Section 2](#).




Figure A2. Distribution of Generative AI Exposure across Broad Occupational Groups

Notes: The figure reports the distribution of the generative AI exposure measure across broad occupational groups, defined at the 2-Digit SOC level, in the OEWS analysis sample. Each box summarizes the distribution of exposure across detailed occupation-by-industry cells within a SOC2 group. Boxes indicate the interquartile range, horizontal lines indicate medians, whiskers show the range excluding outside values, and points denote outliers. Cells are weighted by baseline 2022 employment.




Figure A3. Distribution of Generative AI Exposure across Knowledge-Intensive Service Industries

Notes: The figure reports the distribution of the generative AI exposure measure across knowledge-intensive service industries in the OEWS analysis sample. Industries are defined at the 2-Digit NAICS level and include sectors 51–56. Each box summarizes the distribution of exposure across detailed occupation-by-industry cells within a NAICS2 sector. Boxes indicate the interquartile range, horizontal lines indicate medians, whiskers show the range excluding outside values, and points denote outliers. Cells are weighted by baseline 2022 employment.




Figure A4. Effect of Generative AI Exposure on the Share with GenAI Skills in Revelio

Notes: The figure plots the dynamic difference-in-differences coefficients from equation (3.3), using the share of active workers associated with GenAI-related skills as the dependent variable. The unit of observation is a firm-by-occupation-by-MSA-by-year cell. The coefficients interact the predetermined cell-level GenAI exposure measure, AI_{fom} , defined in Section 2.2, with year indicators and are measured relative to the omitted year 2022. The specification also includes year-specific interactions with pre-ChatGPT occupational work-from-home exposure and firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Observations are weighted by average cell employment during 2018–2022. Standard errors are two-way clustered by firm and occupation, and the shaded region denotes a 95% confidence interval.




Figure A5. Effect of Generative AI Exposure on Posted Wages in Revelio

Notes: The figure plots the dynamic difference-in-differences coefficients using the real log posted wage from Revelio job advertisements as the dependent variable. The coefficients interact the predetermined cell-level GenAI exposure measure, AI_{fom} , defined in Section 2.2, with year indicators and are measured relative to the omitted year 2022. The specification also includes year-specific interactions with pre-ChatGPT occupational work-from-home exposure and firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Observations are weighted by the average number of vacancies in the cell during 2018–2022. Standard errors are two-way clustered by firm and occupation, and the shaded region denotes a 95% confidence interval. The corresponding pooled estimate is reported in column (2) of Table 11.

Table A1. Summary Statistics by Seniority: Revelio Analysis Sample

Variable	Junior	Mid-level	Senior
Real log wage	10.007	10.544	11.035
Employment stock	102.886	62.785	44.380
Separation rate	0.165	0.147	0.120
Job-to-job transition rate	0.094	0.085	0.064
Persistent-non-employment separation rate	0.065	0.055	0.049
New positions	24.053	15.055	8.272
Ending positions	21.481	12.382	7.252
Employment share	0.507	0.285	0.208
New-position share	0.507	0.308	0.185
AI exposure	0.569	0.579	0.572
WFH exposure	0.048	0.048	0.048
GenAI worker share	0.004	0.006	0.007

Notes: The table reports weighted means by seniority group for firm–occupation–MSA–year observations in knowledge-intensive industries, defined as NAICS two-digit sectors 51–56. The sample is a balanced panel covering 2018–2026 and is restricted to cells with more than one employee in 2022. Wages are real log wages for newly created positions and are weighted by the corresponding seniority group’s average number of new positions during 2018–2022. Employment stocks, worker-flow rates, shares, and exposure measures are weighted by the cell’s average employment during 2018–2022. Within each seniority group, the separation rate is defined as annual group separations divided by the corresponding lagged group employment stock. The job-to-job transition rate and persistent-non-employment separation rate are constructed analogously, using group-specific job-to-job transitions and separations without observed re-employment within 90 days as their respective numerators. Observations with missing or nonpositive lagged group employment are excluded from the corresponding rate statistic. Within each seniority group, observations for which any of the three group-specific worker-flow rates exceeds one are excluded from all three rate statistics for that group. This restriction does not affect the remaining seniority-specific descriptive statistics. New and ending positions are counts measured at the firm–occupation–MSA–year level. Employment shares and new-position shares report the fraction of the corresponding cell total accounted for by each seniority group. Junior, mid-level, and senior correspond to Revelio’s seniority categories. AI exposure is the predetermined cell-level measure constructed from the detailed O*NET occupational composition of the firm–occupation–MSA cell, while WFH exposure is defined at the occupation level. Their means are calculated over cells with positive employment in the corresponding seniority group.

Table A2. Generative AI Exposure, GenAI Skills, and Wages in Revelio

	(1) Baseline	(2) No positive GenAI share	(3) Positive GenAI share	(4) GenAI skilled worker (a)	(5) GenAI skilled worker (b)	(6) GenAI skilled worker (c)
Post-ChatGPT × AI exposure	−0.140*** (0.032)	−0.065*** (0.024)	−0.172*** (0.034)	−0.114*** (0.027)		−0.111*** (0.027)
GenAI worker share					−0.151*** (0.020)	−0.122*** (0.017)
GenAI-skilled worker				0.083*** (0.020)	0.089*** (0.021)	0.088*** (0.021)
Post-ChatGPT × WFH exposure	0.422** (0.177)	0.279* (0.154)	0.410* (0.242)	0.391** (0.167)	0.216 (0.165)	0.384** (0.166)
Observations	4,896,686	4,233,318	579,755	8,797,481	8,797,481	8,797,481
R-squared	0.891	0.861	0.919	0.872	0.872	0.872
Firm-by-occupation-by-MSA FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
MSA-by-year FE	Yes	Yes	Yes	Yes	Yes	Yes
WFH control	Yes	Yes	Yes	Yes	Yes	Yes
Sample	2018–2026	2018–2026	2018–2026	2018–2026	2018–2026	2018–2026

Notes: The table reports estimates for the real log starting wage of newly created positions in Revelio. Column (1) presents the baseline pooled difference-in-differences specification. Columns (2) and (3) divide firm–occupation–MSA cells according to whether their GenAI worker share is ever positive during 2018–2026. Column (2) includes cells that never report a positive GenAI worker share, while column (3) includes cells with a positive share in at least one year. Columns (4)–(6) divide newly created positions according to workers’ GenAI-skill status. Column (4) controls for the interaction between predetermined cell-level AI exposure and the post-ChatGPT indicator, column (5) controls for the GenAI worker share in the corresponding firm–occupation–MSA–year cell, and column (6) includes both measures. The GenAI-skilled-worker coefficients therefore measure the conditional starting-wage differential between workers with and without observed GenAI-related skills. All specifications control for the interaction between pre-ChatGPT occupational work-from-home exposure and the post-ChatGPT indicator and include firm-by-occupation-by-MSA, firm-by-year, and MSA-by-year fixed effects. Columns (1)–(3) are weighted by the average number of newly created positions in the corresponding cell during 2018–2022, while columns (4)–(6) are weighted by the average number of newly created positions within each GenAI-skill-status group during the same period. Columns (2)–(6) should be interpreted as descriptive conditional associations because the presence and reporting of GenAI-related skills may respond endogenously to wages, worker composition, and other cell characteristics. Columns (4)–(6) use a common estimation sample with nonmissing information on worker GenAI-skill status, cell-level GenAI worker share, predetermined cell-level AI exposure, and work-from-home exposure. Because re-estimating the fixed effects separately by GenAI-share group may generate group-specific singleton observations, the observation counts in columns (2) and (3) need not sum exactly to the observation count in column (1). Standard errors are two-way clustered by firm and occupation and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

B Harmonization of Occupation and Industry Classifications

A central challenge in constructing a long panel from the OEWS is that both occupational and industry classifications change over time. Occupations are reported under different versions of the Standard Occupational Classification (SOC), while industries are reported under different versions of the North American Industry Classification System (NAICS). Since our empirical strategy compares the evolution of wages and employment within occupation-by-industry cells before and after the release of ChatGPT, these cells must be comparable over time.

Without harmonization, changes in classification systems could mechanically generate spurious entry, exit, or changes in wages and employment. For example, an occupation that is split into two new codes after a classification revision could appear to shrink or disappear, even if the underlying labor market activity did not change. Similarly, if several industry codes are consolidated into a new code, measured employment and wages may change mechanically because of the coding revision rather than because of actual economic adjustment. Harmonization is therefore necessary to ensure that our panel tracks economically comparable occupation-by-industry cells throughout the sample period.

The harmonization procedure has three goals. First, it assigns each occupation and industry code to a consistent classification unit over time. Second, it avoids artificially splitting cells when a historical code maps into multiple new codes. Third, it aggregates cells when several historical codes are consolidated into a single later code. This produces a balanced panel in which changes in outcomes reflect changes in wages and employment rather than changes in the statistical classification system.

B.1 Occupational harmonization

OEWS occupations are reported using the 2010 SOC system in the earlier part of the sample and the 2018 SOC system in the later part of the sample. To construct comparable occupation cells, we use the official SOC crosswalk between the 2010 and 2018 systems. The crosswalk identifies one-to-one mappings, many-to-one mappings, and one-to-many mappings between the two classifications.

We apply a directional harmonization rule. When several 2010 occupations map into a single 2018 occupation, we use the 2018 occupation as the harmonized code. This aggregates the earlier occupations into the later consolidated classification. When one 2010 occupation maps into several 2018 occupations,

we preserve the 2010 occupation as the harmonized code. This avoids artificially splitting a pre-existing occupation into multiple categories that cannot be separately observed in the earlier data. When the mapping is one-to-one, we use the 2018 occupation as the harmonized code.

Formally, let o^{2010} denote an occupation under the 2010 SOC system and o^{2018} the corresponding occupation under the 2018 SOC system. The harmonized occupation o^H is defined as follows:

$$o^H = \begin{cases} o^{2018}, & \text{if several 2010 occupations map into one 2018 occupation,} \\ o^{2010}, & \text{if one 2010 occupation maps into several 2018 occupations,} \\ o^{2018}, & \text{if the mapping is one-to-one.} \end{cases} \quad (\text{A1})$$

We then create a version-specific mapping that assigns each original SOC code in each year to its harmonized occupation. Years reported under the 2010 SOC system are mapped from the 2010 code to the harmonized code, while years reported under the 2018 SOC system are mapped from the 2018 code to the harmonized code. We keep deterministic mappings and drop cases in which an original code cannot be uniquely assigned to a harmonized occupation.

After this step, every occupation in the OEWS panel is expressed in a common harmonized SOC classification. We then merge the generative AI exposure measure to this harmonized detailed occupation code. This ensures that the exposure measure is assigned consistently over time and that post-2018 SOC revisions do not mechanically change treatment status.

B.2 Industry harmonization

The OEWS industry data also span several NAICS revisions. The earlier years use the 2012 NAICS system, the middle years use the 2017 NAICS system, and the later years use the 2022 NAICS system. Since our empirical analysis is conducted at the four-digit industry level, we harmonize NAICS codes at the four-digit level.

The harmonization proceeds in two stages. First, we harmonize the 2012 and 2017 NAICS systems. We begin with the official six-digit crosswalk and collapse it to the four-digit level. We then identify whether each mapping is one-to-one, many-to-one, or one-to-many. If several 2012 industries map into one 2017 industry, we use the 2017 code as the harmonized industry. If one 2012 industry maps into

several 2017 industries, we preserve the 2012 code as the harmonized industry. If the mapping is one-to-one, we use the 2017 code.

Second, we harmonize the 2017 and 2022 NAICS systems using the same logic. If several 2017 industries map into one 2022 industry, we use the 2022 code. If one 2017 industry maps into several 2022 industries, we preserve the 2017 code. If the mapping is one-to-one, we use the 2022 code.

This produces directional harmonized mappings for the 2012–2017 transition and the 2017–2022 transition. We then combine these mappings into a single version-specific harmonization rule for all years. For codes reported under the 2012 system, the code is first mapped through the 2012–2017 bridge and then, when possible, through the final harmonized 2017 classification. For codes reported under the 2017 system, we apply the final 2017 harmonization rule directly. For codes reported under the 2022 system, we map back through the 2017 bridge when necessary and otherwise apply the 2017–2022 rule.

The resulting harmonized industry code is a four-digit NAICS unit that is comparable across all classification regimes. We also construct the corresponding three-digit and two-digit industry groupings from the harmonized four-digit code. These harmonized industry codes are used to define the analysis sample, fixed effects, clusters, and inverse-probability weights.

B.3 Aggregation after harmonization

After harmonizing occupations and industries, some original OEWS cells map into the same harmonized occupation-by-industry cell. In those cases, we aggregate the original cells to the harmonized level. Employment is summed across all original cells that belong to the same harmonized occupation-by-industry-year cell:

$$L_{ojt}^H = \sum_{k \in \mathcal{H}(o,j,t)} L_{kt}, \quad (\text{A2})$$

where $\mathcal{H}(o,j,t)$ denotes the set of original OEWS cells that map into harmonized occupation o and harmonized industry j in year t .

Wage outcomes are aggregated using employment weights:

$$W_{ojt}^H = \frac{\sum_{k \in \mathcal{H}(o,j,t)} L_{kt} W_{kt}}{\sum_{k \in \mathcal{H}(o,j,t)} L_{kt}}. \quad (\text{A3})$$

Thus, when multiple original cells are combined, larger cells receive more weight in the harmonized

wage measure. This aggregation is applied consistently to the wage statistics used in the analysis.

The final OEWS panel is therefore defined at the harmonized detailed occupation by harmonized four-digit industry by year level. This panel is the basis for the OEWS event-study analysis. The harmonization ensures that the identifying variation comes from changes in wages and employment across comparable labor market cells, rather than from changes in the SOC or NAICS classification systems.