

Generative Artificial Intelligence Exposure and U.S. Occupational Wage Polarization: Early Evidence and Workforce-Education Implications from 2018-2025 Occupational Data

Aiden Hanjoon Lee 

Social Science Division/STEM Science Center, 111 Charlotte Place Ste#100/Englewood Cliffs, NJ 07632, USA

Citation: Aiden Hanjoon Lee (2026). Generative Artificial Intelligence Exposure and U.S. Occupational Wage Polarization: Early Evidence and Workforce-Education Implications from 2018-2025 Occupational Data. *Journal of Business, IT, and Social Science*.

DOI: <https://doi.org/10.51470/BITS.2026.05.01.104>

Corresponding Author: **Aiden Hanjoon Lee** | E-Mail: ALee@STEMsc.org

09 February 2026: Received | 12 March 2026: Revised | 11 April 2026: Accepted | 13 May 2026: Available Online

Copyright: This is an open access article distributed under the terms of the Creative Commons Attribution License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

ABSTRACT

Aims: This study examines whether U.S. occupations with greater exposure to large language models showed weaker wage and employment performance during the first measurable years after the public release of ChatGPT, and whether the observed pattern has implications for workforce education and social inequality.

Study design: Longitudinal occupation-level empirical analysis using national U.S. occupational wage and employment data.

Place and Duration of Study: United States national labor market; Bureau of Labor Statistics Occupational Employment and Wage Statistics data for 2018-2025, analyzed in 2026.

Methodology: Annual occupation-level median wages and employment estimates were merged with published large-language-model exposure scores by six-digit Standard Occupational Classification code. Occupations were grouped into four exposure bins, and post-2022 changes in real median wages and employment were compared. A descriptive ordinary least squares trend was estimated between continuous AI exposure and real wage growth.

Results: Employment growth from 2022 to 2025 declined as exposure increased: low-exposure occupations grew by about 5.0%, medium-low exposure occupations by 6.1%, medium-high exposure occupations by 2.6%, and high-exposure occupations by only 0.5%. High-exposure occupations did not receive a broad real wage premium. The occupation-level wage trend was negative, with an estimated slope of approximately -2.9 percentage points from the least-exposed to the most-exposed end of the distribution, although dispersion across occupations remained substantial.

Conclusion: The evidence is best interpreted as an early descriptive signal, not a causal estimate. The results are more consistent with early AI-related wage stagnation and slower job-count growth in highly exposed occupations than with an immediate high-skill wage boom. Workforce education should therefore emphasize AI literacy, domain expertise, statistical reasoning, communication, ethics, and human judgment rather than narrow tool use alone.

Keywords: artificial intelligence; large language models; wage polarization; occupational exposure; workforce education; employment; labor economics; automation.

Abbreviations

AI: artificial intelligence; BLS: Bureau of Labor Statistics; CPI-U: Consumer Price Index for All Urban Consumers; CPS: Current Population Survey; LLM: large language model; OEWS: Occupational Employment and Wage Statistics; SOC: Standard Occupational Classification.

1. Introduction

The public release of widely available generative artificial intelligence tools changed the practical meaning of workplace automation. Earlier digital technologies often affected routine clerical, production, or calculation tasks first. By contrast, large language models can draft text, summarize documents, write and debug code, translate, classify information, search across documents, and support routine administrative communication. These capabilities place a visible part of professional, educational, and office-based work within the reach of automation or computer-assisted redesign [14, 15].

This shift matters for society and education because many occupations once viewed as relatively protected from automation are language-intensive. Writers, translators, public relations specialists, clerical workers, customer service representatives, survey researchers, statisticians, and programmers all perform tasks that can now be accelerated or partly imitated by LLM-based tools. The central question is therefore not whether AI can perform isolated tasks, but whether labor-market outcomes already look different where such tasks are common.

The existing labor economics literature gives competing expectations. One tradition emphasizes skill-biased technological change, in which new technologies raise the relative demand for workers who perform abstract and analytical tasks [1]. A related literature on job polarization shows that employment growth can shift away from routine middle-skill occupations toward high-wage abstract work and low-wage service work [17].

Generative AI does not fit neatly into this older pattern. It can complement workers by lowering search costs, assisting with writing, and spreading good practices. It can also substitute for workers when the final product is standardized text, routine code, correspondence, translation, or simple analysis. Experimental evidence shows that generative AI can raise productivity in some knowledge-work settings, especially for less-experienced workers, but the distribution of those gains between workers, firms, and consumers remains uncertain [22, 7].

For this reason, exposure should not be equated with displacement. A high exposure score means that the task content of an occupation is technically reachable by LLMs. It does not mean that employers have adopted the technology, that output quality is adequate, that regulation allows substitution, or that customers accept machine output. Occupations may be exposed and still resilient if they require accountability, licensing, field knowledge, human trust, or judgment under uncertainty [5, 3].

Official occupational data provide a useful but imperfect first view. The Bureau of Labor Statistics Occupational Employment and Wage Statistics program reports annual national employment and wage estimates for detailed occupational categories in the 2018 Standard Occupational Classification system [9, 10]. These data allow occupational AI-exposure scores to be connected with observed changes in wages and job counts.

The same data requires caution. OEWS is an establishment-based occupational survey. It measures jobs and wage distributions, not complete worker histories. It cannot directly show whether a job is full-time, part-time, contract-based, platform-mediated, or stable across time. This limitation is especially important for AI research because the labor market can change before mass unemployment appears. Firms can reduce hours, rely more on contractors, slow entry-level hiring, or increase workloads while keeping headline occupational job counts relatively stable.

This paper asks a deliberately narrow question: during the first measurable years after the public release of ChatGPT, did occupations with higher LLM exposure show weaker wage and employment trajectories than occupations with lower exposure? The contribution is threefold. First, the study connects published LLM exposure scores with the newest national occupational wage and employment data available for 2018-2025. Second, it distinguishes between complementarity and substitution by comparing low, medium-low, medium-high, and high exposure occupations. Third, it interprets the findings for workforce education, emphasizing the skills that may help workers remain resilient when language and information-processing tasks become partly automatable.

Three hypotheses guide the analysis. First, high-exposure occupations should show weaker post-2022 employment growth if employers can reorganize tasks around AI tools. Second, highly exposed workers may not receive a compensating wage premium if AI reduces the scarcity value of some cognitive tasks. Third, the strongest labor-market outcomes may occur in a middle range of exposure, where AI supports work without substituting for the core task bundle. The analysis is not designed to prove causation. It is an early-warning descriptive analysis of occupational patterns during a period of unusually rapid technology diffusion.

2. Materials and Methods

2.1 Study design and unit of analysis

The study uses a longitudinal ecological design at the detailed-occupation level.

The unit of analysis is the six-digit Standard Occupational Classification occupation. The empirical design compares occupational groups over time rather than following individual workers. The analysis can therefore identify patterns consistent with AI-related pressure, but it cannot prove that AI caused a particular occupational wage or employment change.

The analytic period is divided into a pre-ChatGPT baseline window (2018-2021), an inflection year (2022), and an early post-ChatGPT outcome window (2023-2025). The year 2022 is used as the reference year for post-period changes because the public release of ChatGPT occurred late in that year and created a practical boundary between limited access to LLMs and broad public experimentation with them.

2.2 Wage and employment data

Annual wage and employment estimates were obtained from the BLS OEWS national occupational tables for 2018-2025. The OEWS program publishes estimates by detailed occupation, including employment counts and wage percentiles. The main wage outcome is the annual median wage because it is less sensitive than the mean to extreme values and is easily interpreted as the wage of a typical job within an occupation [9].

Nominal annual median wages were converted to 2022 dollars using the CPI-U. The adjustment is necessary because the period after 2021 included elevated inflation; without deflation, nominal wage increases would overstate changes in purchasing power. Employment was measured as the total estimated number of jobs in occupation i in year t . Occupations with missing SOC codes, suppressed wage estimates, or non-comparable coding should be removed or harmonized before final replication. The May 2025 OEWS estimates use the 2018 SOC and include approximately 830 occupational categories, so the final submission should provide a supplementary SOC crosswalk and exclusion list [10].

2.3 AI exposure scores

Occupation-level AI exposure scores were drawn from a published LLM exposure framework [14]. This exposure measure estimates the share of occupational tasks that could be meaningfully affected by LLMs, particularly when LLMs are combined with software tools. Scores range from 0 to 1, with higher values indicating greater technical exposure of the task bundle to LLM-enabled assistance or automation [14].

The exposure scores were merged with the OEWS panel by six-digit SOC code. In this analysis, exposure is time-invariant. It measures potential task exposure rather than actual adoption by a firm, industry, or region. A high score indicates that the tasks are technologically reachable; it does not establish that workers were replaced or that AI tools were used successfully in the workplace.

2.4 Exposure groups and outcome measures

For descriptive presentation, occupations were grouped into four exposure bins: Low exposure (0.00-0.25), Medium-Low exposure (>0.25-0.50), Medium-High exposure (>0.50-0.75), and High exposure (>0.75-1.00). Binning reduces the influence of outlier occupations and allows broad comparison across the exposure distribution. The main outcomes are real median annual wage by exposure bin, aggregate employment change by exposure bin from 2022 to 2025, occupation-level real wage growth from 2022 to 2025, and wage and employment changes among the 20 detailed occupations with the highest exposure scores.

Table 1: Conceptual interpretation of AI exposure groups used in the analysis

Exposure group	Score range	Typical task profile	Expected labor-market mechanism
Low	0.00-0.25	Physical, location-specific, manual, interpersonal, or embodied tasks.	Limited direct LLM substitution; wage changes mainly reflect local service demand and labor supply.
Medium-Low	>0.25-0.50	Mixed tasks with documentation, communication, and coordination.	Potential complementarity: AI can improve productivity while human judgment remains central.
Medium-High	>0.50-0.75	Substantial text, data, coding, clerical, analytical, or communication tasks.	Ambiguous zone: AI can assist workers but may also reduce demand for routine components.
High	>0.75-1.00	Strong dependence on language, writing, translation, clerical processing, or standardized analysis.	Higher substitution risk and possible wage stagnation if AI reduces the scarcity of occupational skills.

2.5 Statistical analysis

For wage trajectories, the median real annual wage was calculated across occupations in each exposure bin and year. For employment trajectories, employment was summed across occupations within each exposure bin, and the percentage change from 2022 to 2025 was calculated as $EmploymentChange_b = 100 \times (Employment_b,2025 - Employment_b,2022) / Employment_b,2022$. Occupation-level real wage growth was calculated in the same way using real median wages.

A simple ordinary least squares model was used to describe the relationship between continuous AI exposure and occupation-level real wage growth from 2022 to 2025: $WageGrowth_i = \alpha + \beta \times AIExposure_i + \epsilon_i$. The coefficient beta is interpreted as the percentage-point change in wage growth associated with a one-unit increase in AI exposure. This regression is descriptive. It is not a causal model because it does not isolate AI adoption from post-pandemic demand shifts, interest-rate effects, sectoral cycles, and other changes.

2.6 Methodological limitation: jobs, workers, hours, and employment quality

The most important data limitation is that OEWS employment estimates count jobs rather than unique workers. A 40-hour permanent job and a lower-hour job both appear as jobs, and multiple jobholding cannot be observed. OEWS also does not directly measure contract status, platform work, benefits, hours, or job stability. This limitation is central to the interpretation of AI labor-market effects. If firms use AI to reduce hours, rely on contractors, or slow entry-level hiring, aggregate job counts could remain stable while employment quality worsens. A stronger follow-up study should link exposure scores to Current Population Survey microdata, American Community Survey data, online vacancy postings, and firm-level AI adoption measures [11, 20].

3. Results

3.1 Distribution and interpretation of occupational AI exposure

The exposure distribution shows that generative AI is not limited to a narrow group of technology occupations. High-exposure occupations include writers, translators, survey researchers, clerical workers, public relations specialists, programmers, statisticians, and other occupations in which language, documentation, information retrieval, coding, or standardized analysis is central to daily work. Low-exposure occupations are more likely to require physical presence, manual manipulation, environmental interaction, personal care, transportation, repair, or protective service.

This pattern explains why the AI transition cannot be reduced to a simple high-skill versus low-skill comparison. Some highly exposed occupations require college or graduate training, while others are clerical or administrative.

Some low-exposure occupations are low wage, but others require substantial skill or certification. The more relevant distinction is the structure of tasks: LLMs perform text, code, communication, and information-processing tasks more directly than site-specific physical work.

3.2 Real wage trajectories by AI exposure bin, 2018-2025

Figure 1 shows real median annual wages by exposure bin from 2018 through 2025, with 2022 marked as the ChatGPT inflection year. The main point is not a sudden collapse in high-exposure wages. Rather, the high-exposure group shows no broad compensating wage premium after 2022. High-exposure occupations begin the period near the upper-middle portion of the wage distribution, but by 2025, they remain below the stronger-performing medium-low and medium-high groups.

The medium-low exposure group shows the strongest wage performance, reaching approximately \$72,000 by 2025. This pattern is consistent with complementarity. In these occupations, AI may reduce documentation burden, improve search and communication, or support analysis, while the core work still depends on human judgment, coordination, domain knowledge, and accountability.

The low-exposure group remains comparatively flat, with real wages in the low-\$40,000 range. This is important because protection from LLM substitution does not automatically produce wage growth. Many manual and service occupations remain constrained by local demand, weak bargaining power, and limited measured productivity growth.

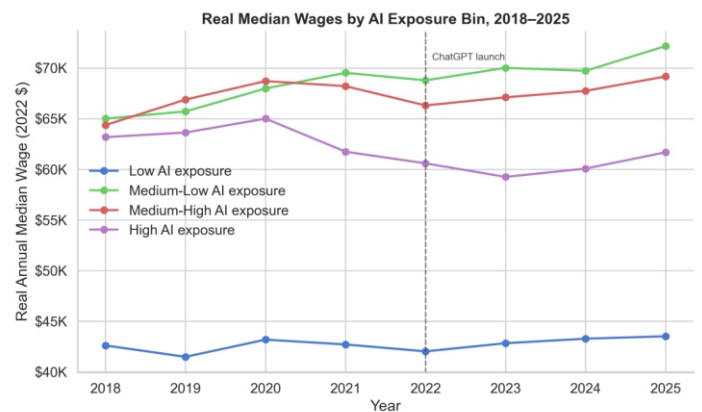


Figure 1: Real median annual wages by AI exposure bin, 2018-2025, expressed in 2022 dollars. The dashed vertical line marks the late-2022 public launch of ChatGPT.

3.3 Employment change by AI exposure bin, 2022-2025

Figure 2 shows aggregate employment change from 2022 to 2025 by AI exposure bin. Employment growth declines as exposure rises: low-exposure occupations grew by approximately 5.0%, medium-low exposure occupations by 6.1%, medium-high exposure occupations by 2.6%, and high-exposure occupations by only 0.5%. This is one of the clearest descriptive findings in the analysis.

The medium-low group again performs best, combining strong wage growth with the highest employment growth. That joint pattern supports the complementarity interpretation. If AI tools help moderately exposed workers handle communication, documentation, scheduling, and routine analysis more efficiently, employers may expand these roles instead of reducing them.

Near-zero growth in high-exposure occupations is meaningful because it occurred during a period in which the broader labor market was not uniformly contracting. At the same time, the result should not be described as proof of displacement. OEWS job counts do not show hours, contract status, multiple jobholding, or job quality. The cautious conclusion is that official job-count growth in high-exposure occupations was markedly weaker than in lower-exposure groups.

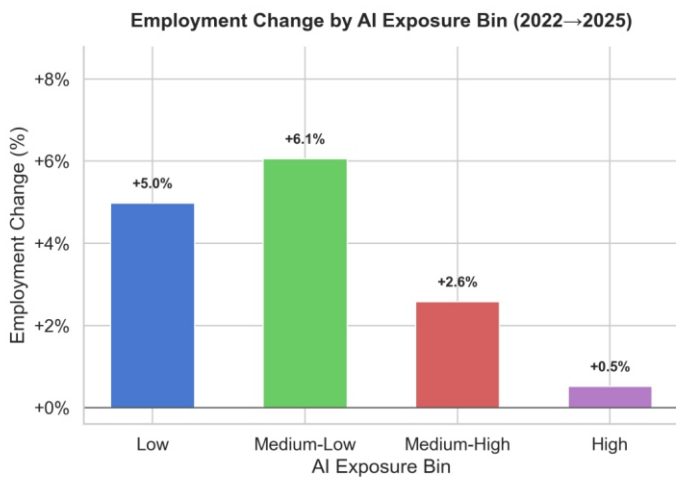


Figure 2: Aggregate employment change by AI exposure bin from 2022 to 2025. Employment growth declines as occupational AI exposure rises.

3.4 Occupation-level association between AI exposure and real wage growth

Figure 3 plots occupation-level real wage growth from 2022 to 2025 against AI exposure. The ordinary least squares trend has a slope of approximately -2.9 percentage points. Interpreted descriptively, moving from the least-exposed to the most-exposed end of the occupational distribution is associated with about 2.9 percentage points lower real wage growth over the post-2022 window.

The scatter also shows wide dispersion. Some low-exposure occupations experienced weak or negative wage growth, and some higher-exposure occupations experienced positive growth. This dispersion is not a weakness of the finding; it is the correct interpretation. AI exposure is one factor among many. Occupational outcomes also reflect industry demand, licensing, geography, education requirements, unionization, remote-work restructuring, and firm-level adoption choices.

Computer and mathematical occupations illustrate this heterogeneity. Some coding tasks are directly affected by LLMs because models can generate, debug, explain, and refactor code.

Yet other technical roles may become more valuable when workers use AI to handle routine coding while focusing on system design, security, data architecture, and domain-specific problem-solving.

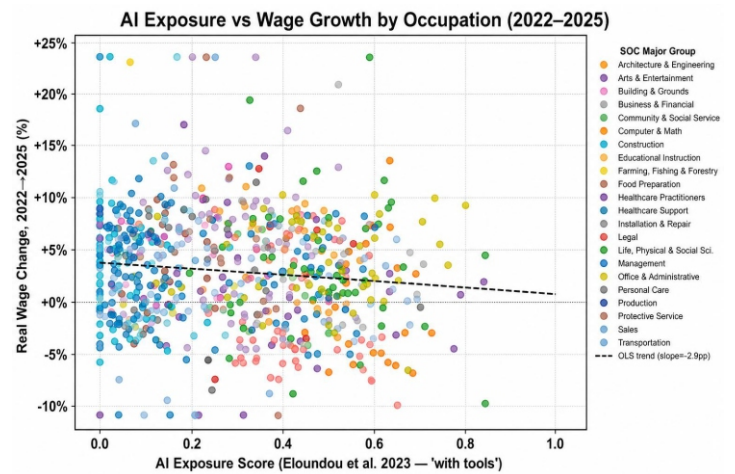


Figure 3: Occupation-level AI exposure score versus real wage growth from 2022 to 2025. Points are colored by SOC major group, and the dashed line shows the ordinary least squares trend.

3.5 Highest-exposure occupations

Table 2 translates the exposure-bin analysis into recognizable occupations. The table also shows why exposure should not be interpreted mechanically. Some high-exposure occupations lost both wages and employment, some lost employment while wages rose, and some gained employment despite high measured exposure.

The most concerning joint pattern appears in computer programmers, where real wages declined by 6.7%, and employment declined by 30.5%. Writers and authors show a similar direction, with wages declining by 4.4% and employment declining by 11.2%. These occupations are directly relevant to the generative-AI transition because code generation and text generation are among the most visible capabilities of current models.

Other occupations show a different pattern. Brokerage clerks, correspondence clerks, advertising sales agents, and customer service representatives experienced employment declines while wages were flat or positive. This may reflect survivor effects: if routine positions shrink, the remaining workers may be more experienced or assigned to higher-value tasks. It may also reflect restructuring in which routine components are automated while human-intensive components remain.

Employment gains among animal scientists, environmental scientists, concierges, and tax examiners warn against a deterministic reading of exposure scores. Some occupations may have language or information-processing tasks but still depend on fieldwork, regulatory knowledge, institutional context, or personal trust. These features can buffer workers from direct substitution.

Table 2: Top 20 highest AI-exposure occupations with real wage and employment changes from 2022 to 2025

Occupation	AI exposure	Wage change, 2022-2025	Employment change, 2022-2025
Animal Scientists	0.844	-9.7%	+23.0%
Survey Researchers	0.844	+4.5%	+5.2%
Interpreters and Translators	0.840	+2.0%	-0.2%
Brokerage Clerks	0.800	+9.3%	-15.8%
Public Relations Specialists	0.788	+0.8%	+7.0%
Writers and Authors	0.774	-4.4%	-11.2%
Legal Secretaries and Administrative Assistants	0.761	+3.6%	-2.3%
Executive Secretaries and Executive Administrative Assistants	0.743	+5.5%	-3.2%
Correspondence Clerks	0.732	+9.9%	-13.7%
Statisticians	0.725	-2.9%	-5.7%
Environmental Scientists and Specialists	0.710	-2.3%	+15.5%
Sales Representatives, Wholesale and Manufacturing	0.707	+3.6%	-2.8%
Customer Service Representatives	0.705	+7.7%	-9.9%
Concierges	0.700	-0.4%	+31.0%
Credit Counselors	0.693	+0.3%	-4.5%
Advertising Sales Agents	0.692	+0.8%	-13.9%
Mathematicians	0.690	+2.7%	-1.9%
Computer Programmers	0.683	-6.7%	-30.5%
Tax Examiners and Collectors, and Revenue Agents	0.677	-2.2%	+11.9%
Counter and Rental Clerks	0.677	+4.8%	+8.1%

Note: Green shading indicates positive change and red shading indicates negative change. Values are descriptive estimates from the merged occupation-level analytic file.

4. Discussion

4.1 Main interpretation

The main pattern is closer to early wage stagnation in highly exposed occupations than to classical wage polarization. In the older polarization story, high-skill abstract workers tend to benefit while routine workers lose ground. The evidence here does not show a broad high-exposure wage premium. Instead, high-exposure occupations have weak employment growth and limited real wage momentum after 2022.

This interpretation is consistent with recent work emphasizing that AI can have both productivity and displacement effects. When AI helps workers accomplish more and creates new tasks, labor demand can rise. When AI performs tasks that workers previously supplied, demand for some labor inputs can fall or wage growth can be suppressed [2, 12, 19].

4.2 Complementarity, substitution, and occupational resilience

The complementarity result fits field evidence showing that AI assistance can improve output and quality in certain knowledge-work settings, but that performance depends on whether the task is inside or outside the technology's useful frontier [13]. The medium-low exposure group may occupy this favorable zone: AI is useful enough to raise productivity, but not close enough to the core job to remove the need for human judgment.

The policy implication is not that high-exposure jobs will disappear. A better reading is that high exposure identifies occupations where task redesign is likely. Whether redesign helps or harms workers depends on training, bargaining power, labor-market institutions, firm choices, and whether productivity gains are shared. Recent calls for pro-worker AI therefore emphasize steering the technology toward capability expansion rather than narrow labor-cost reduction [4].

4.3 Workforce-education and social implications

The findings have direct relevance for education, training, and social policy. If high exposure is associated with weaker wage and employment momentum, then simply teaching students or workers to use the latest AI interface is not enough.

Training should help learners combine AI fluency with durable human capabilities: domain knowledge, statistical reasoning, writing and communication, ethical judgment, source evaluation, accountability, and problem framing.

The results also suggest that workforce programs should not treat all white-collar skills as equally safe. Language, coding, clerical processing, translation, and standardized analysis are important, but they may become less scarce when AI systems produce adequate first drafts at low cost. Resilience may depend on moving from routine production of text or code toward verification, integration, client communication, interdisciplinary judgment, and responsibility for consequences.

For educational institutions, the practical message is to redesign curricula around human-AI collaboration rather than AI avoidance. Students should learn when AI output is useful, when it fails, how to audit errors, how to cite and disclose AI assistance, and how to connect AI-generated information to real-world evidence. Such training is especially important for young workers entering occupations where entry-level tasks are most exposed.

4.4 Why the findings remain preliminary

The time window is short. Many organizations were still experimenting with generative AI during 2023-2025, and adoption differs across firms, regions, industries, and occupations. Some effects may appear only after software systems become more integrated into business processes. Conversely, some early effects may reflect post-pandemic adjustment, macroeconomic policy, or industry-specific cycles rather than AI alone.

For that reason, this paper should be read as a technically grounded descriptive warning, not a final causal claim. The pattern is strong enough to justify additional research and workforce attention, but not strong enough to identify the exact mechanism in each occupation.

5. Limitations

Causal identification: The study is descriptive and cannot prove that AI caused observed wage or employment changes.

Other post-2022 forces, including inflation, interest rates, post-pandemic restructuring, and sector-specific demand shocks, may contribute to the patterns.

Short observation window: The period from 2022 to 2025 captures only the earliest phase of generative-AI diffusion. Full labor-market effects may require several additional years to appear.

Occupation-level aggregation: The analysis uses occupations rather than individual workers. It cannot observe age, education, experience, race, gender, hours, firm, region, or individual AI use.

Employment quality: OEWS counts jobs and does not directly measure full-time status, contract status, benefits, income volatility, hours, or multiple jobholding.

Exposure-score limits: LLM exposure scores measure potential task exposure, not actual adoption, quality of AI output, employer strategy, or worker adaptation.

SOC harmonization: Changes in occupational definitions or data suppression can affect longitudinal comparisons unless carefully harmonized.

Weighting and sensitivity: A final replication package should report both occupation-weighted and employment-weighted results because changes in large occupations affect more workers than changes in small occupations.

Regional and firm-level adoption: National averages may hide geographic and firm-level heterogeneity in AI adoption, remote work, and occupational restructuring.

6. Conclusions

This study provides early descriptive evidence that occupational exposure to generative AI is associated with weaker post-2022 labor-market momentum in the United States. Employment growth declined as AI exposure increased, and high-exposure occupations did not show a broad real wage premium. The strongest performance appeared in the medium-low exposure group, suggesting that workers may benefit most when AI complements rather than substitutes for core occupational tasks.

The results should be read cautiously. The post-ChatGPT window is short, and OEWS data cannot measure worker-level hours, contract status, benefits, or job stability. Even so, the patterns are important because labor-market disruption may first appear as slower growth, weaker wage momentum, reduced entry-level hiring, and poorer employment quality rather than sudden visible unemployment.

For education and workforce policy, the practical message is straightforward. AI literacy should be paired with domain knowledge, statistical reasoning, communication, ethics, and human judgment. Workers are likely to be more resilient when they can use AI to expand their productivity while retaining responsibility for tasks that require trust, context, accountability, and real-world judgment. Future research should test these patterns with CPS microdata, online vacancy data, regional variation, and firm-level adoption measures.

Data availability: The wage and employment data are available from the U.S. Bureau of Labor Statistics Occupational Employment and Wage Statistics program. AI exposure scores are based on published occupation-level LLM exposure measures. A final replication package should include the merged analytic dataset, SOC crosswalk, cleaning code, and figure-generation scripts as supplementary materials or in a public repository.

Ethical approval and consent: Not applicable. This study uses public aggregate labor-market data and does not involve identifiable human participants, private personal information, intervention, or identifiable individual records.

Funding: No external funding was received for this study.

Competing interests: The author declares that no competing interests exist.

Acknowledgments: Acknowledgments will be provided after peer review to preserve double-blind review.

Declarations

References

1. Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4B, pp. 1043-1171). Elsevier. [https://doi.org/10.1016/S0169-7218\(11\)02410-5](https://doi.org/10.1016/S0169-7218(11)02410-5)
2. Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488-1542. <https://doi.org/10.1257/aer.20160696>
3. Acemoglu, D., & Restrepo, P. (2019). Artificial intelligence, automation, and work. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 197-236). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.003.0008>
4. Acemoglu, D., Autor, D., & Johnson, S. (2026). Building pro-worker artificial intelligence. *NBER Working Paper No. 34854*. <https://doi.org/10.3386/w34854>
5. Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3-30. <https://doi.org/10.1257/jep.29.3.3>
6. Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279-1333. <https://doi.org/10.1162/003355303322552801>
7. Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889-942. <https://doi.org/10.1093/qje/qjae044>
8. Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108, 43-47. <https://doi.org/10.1257/pandp.20181019>
9. Bureau of Labor Statistics. (2026a). Occupational Employment and Wage Statistics tables: May 2025 national occupational employment and wage estimates. U.S. Department of Labor. <https://www.bls.gov/oes/tables.htm>
10. Bureau of Labor Statistics. (2026b). Technical notes for May 2025 OEWS estimates. U.S. Department of Labor. https://www.bls.gov/oes/current/oes_tec.htm
11. Bureau of Labor Statistics. (2026c). Current Population Survey: Concepts and definitions. U.S. Department of Labor. <https://www.bls.gov/cps/definitions.htm>

12. Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E., & Tavares, M. M. (2024). Gen-AI: Artificial intelligence and the future of work. *IMF Staff Discussion Note* SDN/2024/001. <https://doi.org/10.5089/9798400262548.006>
13. Dell'Acqua, F., McFowland, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., et al. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Management Working Paper No. 24-013*. <https://www.hbs.edu/faculty/Pages/item.aspx?num=64700>
14. Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306-1308. <https://doi.org/10.1126/science.adj0998>
15. Felten, E. W., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195-2217. <https://doi.org/10.1002/smj.3286>
16. Felten, E. W., Raj, M., & Seamans, R. (2023). How will language modelers like ChatGPT affect occupations and industries? *arXiv:2303.01157*. <https://arxiv.org/abs/2303.01157>
17. Goos, M., Manning, A., & Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8), 2509-2526. <https://doi.org/10.1257/aer.104.8.2509>
18. Hampole, M., Papanikolaou, D., Schmidt, L. D. W., & Seegmiller, B. (2025). Artificial intelligence and the labor market. *NBER Working Paper No. 33509*. <https://doi.org/10.3386/w34854>
19. Huang, Y. (2024). The labor market impact of artificial intelligence: Evidence from US regions. *IMF Working Paper WP/24/199*. International Monetary Fund. <https://doi.org/10.5089/9798400287916.001>
20. Katz, L. F., & Krueger, A. B. (2019). The rise and nature of alternative work arrangements in the United States, 1995-2015. *ILR Review*, 72(2), 382-416. <https://doi.org/10.1177/0019793918820008>
21. Machovec, C. (2025). Incorporating AI impacts in BLS employment projections. *Monthly Labor Review*. U.S. Bureau of Labor Statistics. <https://www.bls.gov/opub/mlr/>
22. Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187-192. <https://doi.org/10.1257/science.adh2586>
23. OpenAI. (2022). Introducing ChatGPT. *OpenAI Blog*. <https://openai.com/blog/chatgpt>
24. Stanford Institute for Human-Centered Artificial Intelligence. (2025). *Artificial Intelligence Index Report 2025*. Stanford University. <https://aiindex.stanford.edu/report/>
25. World Economic Forum. (2025). *The Future of Jobs Report 2025*. World Economic Forum. <https://www.weforum.org/reports/the-future-of-jobs-report-2025/>