Do AI Exposure Measures Predict Labor-Market Outcomes? Evidence from Unemployment, Employment and Wages, 2022-2026

Benazir Rowe, PhD
Opportunity Data
benazir@opportunitydata.org
Working paper, November 2026

Abstract

Occupational measures of exposure to artificial intelligence are now used to steer training funds, program review and labor-market forecasts, but they are rarely tested against outcomes. This paper tests five measures against what happened to US occupations after November 2022: the Opportunity Data AI Exposure Index, which scores the structure of work rather than AI capability; the Felten, Raj and Seamans language-modeling AIOE; the Eloundou et al. GPT-4 and human exposure ratings; and the Bureau of Labor Statistics AI Exposure Categories. Using harmonized occupation groups from the Current Population Survey (2015 to September 2026) and the Occupational Employment and Wage Statistics (2016 to 2025), within occupational groups and with placebo windows before generative AI, I find that the Opportunity Data AI Exposure Index is the strongest predictor of the post-2022 rise in unemployment (+0.10 percentage points per year per standard deviation, p=0.019; +4.3% per year as a proportional change; shuffle-test p=0.005). Felten et al. also predicts the proportional rise; the other measures do not reach significance. Holding remote-work potential constant, only the Opportunity Data AI Exposure Index still predicts the rise. All five measures detect the cross-sector decline in young workers' employment equally; none predicts employment or wage change once sector composition is controlled for. Scores are frozen and a forward test is registered for 2027 and 2028.

Keywords: artificial intelligence, occupational exposure, unemployment, validation, O*NET, Current Population Survey

JEL codes: J21, J23, J24, J64, O33

Data and code: a replication package reproducing every table and figure is available at https://opportunitydata.org/ai-exposure-outcome-validation.html.

1. Introduction

Measures of occupational exposure to AI have multiplied since 2021. Most rate what AI can do against the tasks or abilities an occupation requires (Felten, Raj and Seamans 2021, 2023; Eloundou et al. 2024); others record what people use AI for (Tomlinson et al. 2025; Handa et al. 2025); the Bureau of Labor Statistics combined five such sources into its first federal AI exposure categories in August 2026 (BLS 2026). These measures now inform decisions with money attached, including which training

programs states fund. Evidence on whether they predict labor-market outcomes is mixed. Payroll data put employment of workers aged 22 to 25 in highly exposed occupations about 19 percent below its counterfactual by June 2026, operating through hiring (Brynjolfsson, Chandar and Chen 2026); survey and administrative data show the same direction for young workers (Atkinson and Yamco 2026; Tucker 2026). Analyses of unemployment and the occupational mix find little aggregate disruption (Gimbel et al. 2025; Massenkoff and McCrory 2026; Fairlie and Wu 2026), and unemployment risk in exposed occupations began rising in early 2022, before November 2022 (Frank et al. 2026). Earlier comparisons found that individual exposure measures predicted unemployment risk poorly (Frank, Ahn and Moro 2023) and that most showed no differential unemployment rise after 2022 (Eckhardt and Goldschlag 2025). This paper adds within-occupational-group estimates, placebo windows before generative AI, controls for remote-work potential, and a registered forward test.

A specific reason to expect capability ratings to perform poorly is vintage. A rating of what GPT-4 could do in March 2023 is a record of GPT-4 in March 2023. By spring 2026 frontier systems operated computers and completed multi-step work end to end, and every capability input to the existing measures predates that change (Rowe 2026). Reviews of the leading capability ratings note the same temporal limit (Lund et al. 2026). If capability is approaching sufficiency for digitally mediated work, ratings of capability should lose their power to separate digital occupations from one another, and what should matter instead is what still binds: face-to-face contact and physical presence. The Opportunity Data AI Exposure Index was built on that premise.

This paper makes three contributions. First, it tests five measures against four outcomes on public data, within occupational groups, with placebo windows that predate generative AI, so that a measure is credited only for predicting change that was not already under way. Second, it tests the vintage argument directly: if capability ratings have saturated, they should fail inside digital work even though they still vary there. Third, it freezes the scores and registers a forward test with fixed dates, so that the claim can fail in public.

The results favor structural measurement on one outcome and favor no measure on the others. Within occupational groups, unemployment rose faster after 2022 in occupations the Opportunity Data AI Exposure Index rates more exposed, on both the points and the proportional scale, with no positive relationship before 2020; Felten et al. predicts the proportional rise at a smaller size, and once remote-work potential is held constant only the Opportunity Data AI Exposure Index does. Inside the more digital half of the workforce the capability measures still rank occupations differently but predict little. Every measure detects the shift away from young workers across sectors, and none separates occupations within a sector. No measure predicts employment or wages once the 2022 to 2024 interest-rate cycle is accounted for.

2. Measures

Table 1 lists the five measures and the AI each describes. The Opportunity Data AI Exposure Index is the composite of three dimensions built from nine O*NET work-context and work-activity variables: digital intensity raises exposure, and human contact and physical anchoring lower it (Rowe 2026). No AI capability or usage data enter it. The other measures rate AI capability (Felten et al.; Eloundou et al.) or combine capability with observed usage (BLS).

Table 1. Measures tested

Measure	What it rates AI as	Published	Built from
Felten, Raj and Seamans LM-AIOE	AI application benchmarks of the 2010s, reweighted for language modeling	2023	Abilities x AI applications
Eloundou, Manning, Mishkin and Rock	GPT-4 and human raters, March 2023	2023 (Science 2024)	Task-level capability ratings
BLS AI Exposure Categories	Three studies from 2020-23 plus Claude and Copilot usage through January 2025	August 2026	Synthesis of capability and usage
Opportunity Data AI Exposure Index	Assumes frontier capability as of April 2026; scores what still binds	September 2026	O*NET work structure, three dimensions

3. Data and design

Outcomes. From the Current Population Survey (CPS) basic monthly files I measure, for each occupation group, the change in the unemployment rate by occupation of last job, log employment growth, and the young-worker gap: log employment growth at ages 22 to 27 minus that at ages 35 to 54 in the same occupation. The base window is December 2021 to November 2022 and the end window November 2025 to September 2026. From the Occupational Employment and Wage Statistics (OEWS) I measure log employment and log median wage growth from May 2022 to May 2025. All changes are expressed per year.

Placebo. Every outcome is recomputed for a window before generative AI: December 2015 to November 2016 against December 2018 to November 2019 in the CPS, and May 2016 to May 2019 in the OEWS. A measure that predicts the placebo as well as the post-2022 window is tracking a pre-existing trend.

Occupations. Codes changed in 2018 (SOC) and 2020 (Census). I pool codes linked by any split or merge in the official crosswalks into occupation groups, giving 738 OEWS groups and 525 CPS groups observed in every window (Appendix B).

Estimation. For occupation group *o* in major group *g*, I estimate

$$\Delta y_o = \alpha_g + \beta z_o + X_o'\gamma + \varepsilon_o$$

where z_o is the exposure measure standardized to an employment-weighted mean of zero and standard deviation of one and signed so that higher means more exposed, α_g are two-digit SOC major group effects, and X_o are controls: the group's own placebo-window change, its base wage, the 2019 to 2022 COVID dip, and the share of its employment in information, finance, professional services and management of companies. Regressions are weighted by base employment with heteroskedasticity-robust standard errors. The major group effects mean that every comparison is office work against office work and health care against health care, so a sector-wide boom or bust does not drive the estimates.

4. Results

4.1 Unemployment

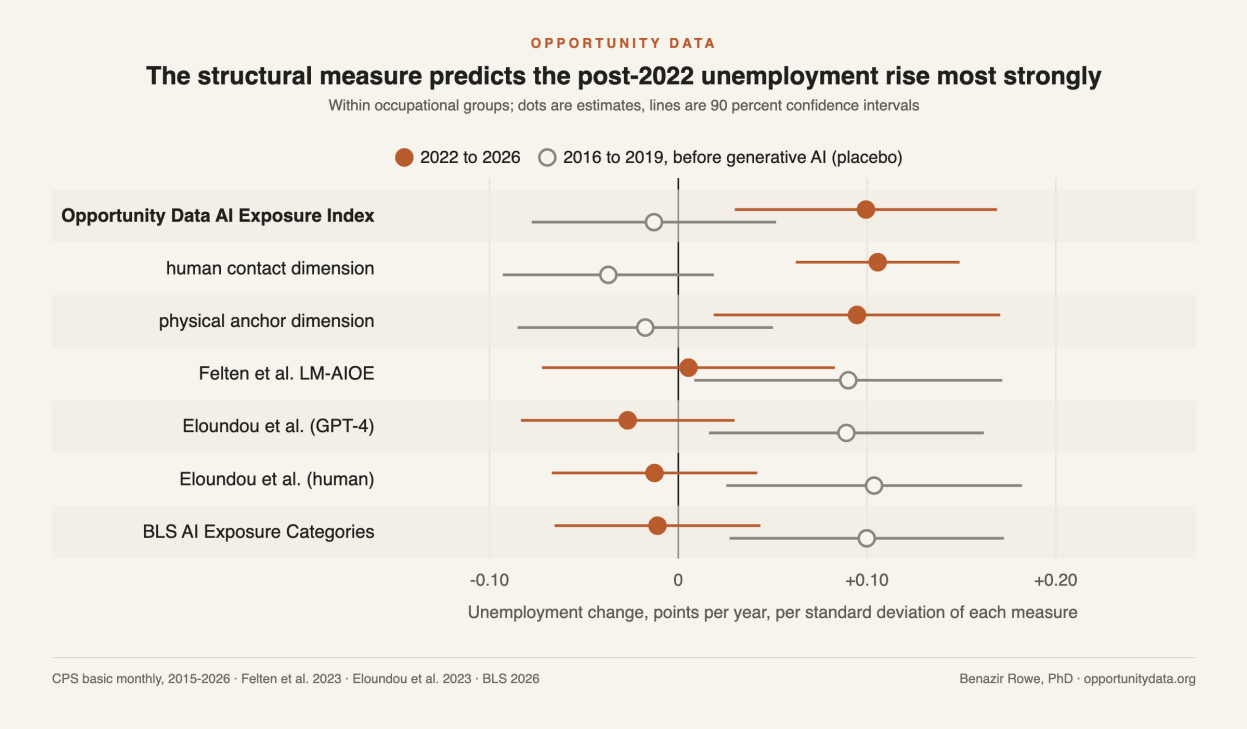
Table 2 reports the main estimates. Within occupational groups, a one standard deviation higher Opportunity Data AI Exposure Index score predicts unemployment rising +0.10 percentage points a year faster (p=0.019), about 0.4 points over the 3.9 years observed. In the placebo window the estimate is -0.01 (p=0.742). On the points scale the four other measures show estimates of -0.03 to +0.01 after 2022, none significant, and +0.09 to +0.10 in the placebo window (p 0.02 to 0.07); Figure 1 shows both windows. Because occupations those measures rate exposed start from lower unemployment, the points scale can understate their post-2022 change; Section 4.3 repeats the test as a proportional change, where their placebo relationship disappears (-1.3% to +1.1%) and Felten et al. predicts the rise (+3.4%, p=0.032). Unemployment risk in exposed occupations began rising before November 2022 (Frank et al. 2026), which is why every estimate here is set against a pre-2020 placebo.

Table 2. Main estimates: outcome change per standard deviation of each measure

Outcome	Opportunity Data AI Exposure Index	Felten et al. LM-AIOE	Eloundou et al. GPT-4	Eloundou et al. human	BLS categories
Unemployment rate, 2022-26	+0.10**	+0.01	-0.03	-0.01	-0.01
	(0.04)	(0.05)	(0.03)	(0.03)	(0.03)
Placebo: unemployment, 2016-19	-0.01	+0.09*	+0.09**	+0.10**	+0.10**
	(0.04)	(0.05)	(0.04)	(0.05)	(0.04)
Ages 22-27 minus 35-54, 2022-26 (across groups)	-1.64***	-1.79***	-1.75***	-1.76***	-1.63***
	(0.48)	(0.53)	(0.41)	(0.50)	(0.45)
Ages 22-27 minus 35-54, 2022-26 (within groups)	-0.16	-0.68	-0.86	-0.53	-0.25
	(1.04)	(1.05)	(0.89)	(1.07)	(0.90)
Employment, OEWS 2022-25 (with sector control)	-0.20	+0.81*	-0.07	+0.29	+0.43
	(0.32)	(0.49)	(0.34)	(0.43)	(0.33)
Median wage, OEWS 2022-25	+0.24	+0.42*	+0.36*	+0.52***	+0.37*
	(0.20)	(0.25)	(0.19)	(0.18)	(0.19)

Notes: Coefficient in percentage points per year (unemployment) or percentage points of growth per year (employment, wages, young-worker gap); robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1. Unemployment rows include major group effects; the young-worker rows are without and with them; employment and wages include major group effects and all controls. 457 to 524 CPS groups and up to 738 OEWS groups. Full results for every outcome and specification are in Appendix D.

Figure 1. Unemployment change per standard deviation, after 2022 and before generative AI



Notes: Estimates with 90 percent confidence intervals, within occupational major groups. The human contact and physical anchor rows are dimensions of the Opportunity Data AI Exposure Index, signed so that higher means more exposed.

4.2 Head to head

Because the measures are correlated (rank correlations of 0.71 to 0.76 between the Opportunity Data AI Exposure Index and each other measure), Table 3 enters the Opportunity Data AI Exposure Index and each competitor in the same regression. The Opportunity Data AI Exposure Index remains at +0.14 to +0.18 (all $p<0.01$); each competitor turns negative (-0.13 to -0.09). The negative competitor coefficients identify occupations rated exposed by capability but anchored by face-to-face work, the same occupations as the divergent set in Section 4.5, so the two results are one fact seen twice, not independent confirmations.

Table 3. Unemployment, 2022-26: the Opportunity Data AI Exposure Index and one competitor in the same regression

Competitor	Opportunity Data AI Exposure Index	Competitor	Occupation groups
Felten, Raj and Seamans LM-AIOE	+0.16 (0.05)	-0.10 (0.06)	353
Eloundou et al., GPT-4 ratings	+0.18 (0.06)	-0.13 (0.06)	353
Eloundou et al., human ratings	+0.14 (0.05)	-0.09 (0.05)	353
BLS AI Exposure Categories	+0.17 (0.05)	-0.12 (0.05)	353

Notes: Percentage points per year per standard deviation, robust standard errors in parentheses. Major group effects, the placebo-window change and the COVID dip are controlled.

4.3 Scale, remote work and chance

Three objections apply to any unemployment result of this kind. First, scale: occupations differ in their starting unemployment rates. As a proportional change, the Opportunity Data AI Exposure Index predicts the rise at +4.3% per year per standard deviation (p=0.006), Felten et al. at +3.4% (p=0.032), and the other three measures at +1.3% to +2.2% (p 0.10 to 0.26); in the 2016-19 placebo its proportional relationship was negative (-2.7%), so the post-2022 rise reverses a pre-AI pattern rather than continuing one. Second, remote work: occupations with little face-to-face contact are also those that can be done from home (Dingel and Neiman 2020), and remote work, return-to-office policies and technology retrenchment all moved after 2022; recent graduates' unemployment in 2026 is linked to remote-work availability (Fairlie and Wu 2026). Holding teleworkability constant, the Opportunity Data AI Exposure Index still predicts the rise, +0.11 points a year (p=0.021) and +3.9% proportionally (p=0.023), while no other measure reaches significance on either scale (proportional estimates +0.4% to +1.9%, all p>0.18). Third, chance: shuffling each measure's scores within occupational groups 1,000 times gives the Opportunity Data AI Exposure Index p=0.005 on the points scale and p<0.001 proportionally, against p 0.39 to 0.86 for the other measures on points and p=0.011 for Felten et al. proportionally (Table 4).

One control weakens the composite. Adding each occupation's late-2022 unemployment rate reduces the Opportunity Data AI Exposure Index to +0.06 (p=0.174), while the human contact dimension holds at +0.11 (p=0.002). That base rate is estimated from small cells and also enters the outcome, so the control is prone to regression to the mean; controlling instead for the 2019 rate, measured before the window, raises the estimate to +0.13 (p=0.003).

Table 4. Unemployment, 2022-26: scale, remote work, starting levels and common sample

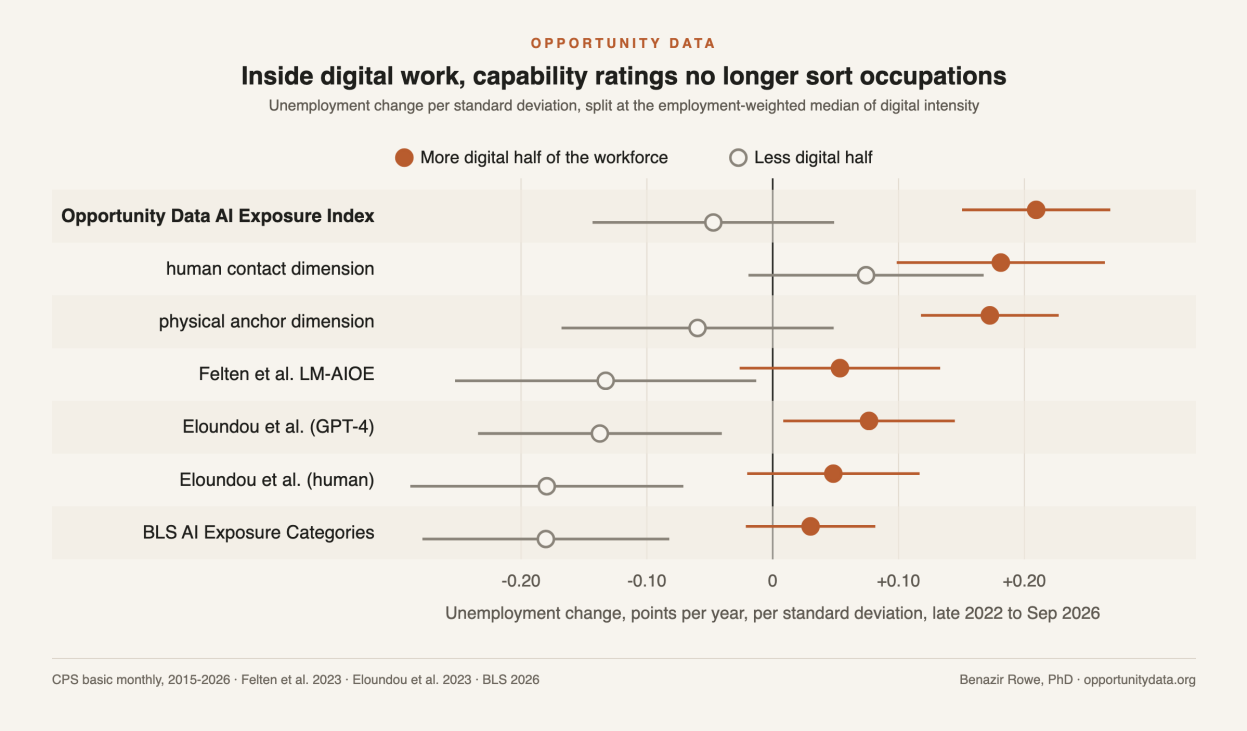
Specification	Opportunity Data AI Exposure Index	Felten et al. LM-AIOE	Eloundou et al. GPT-4	Eloundou et al. human	BLS categories	Remote-work potential
Points (main)	+0.110**	+0.007	-0.032	-0.010	-0.016	+0.056
	(0.045)	(0.050)	(0.039)	(0.037)	(0.037)	(0.037)
Points, remote work held constant	+0.113**	-0.038	-0.057	-0.027	-0.039	
	(0.049)	(0.051)	(0.041)	(0.038)	(0.035)	
Points, 2019 unemployment rate controlled	+0.128***	+0.025	-0.012	+0.003	-0.002	+0.064*
	(0.043)	(0.050)	(0.042)	(0.036)	(0.036)	(0.036)
Points, late-2022 unemployment rate controlled	+0.062	-0.063	-0.135***	-0.052	-0.079**	+0.026
	(0.045)	(0.055)	(0.048)	(0.044)	(0.039)	(0.037)
Points, common sample	+0.110**	+0.019	-0.023	-0.014	-0.017	+0.049
	(0.045)	(0.054)	(0.045)	(0.042)	(0.041)	(0.041)
Proportional	+0.043***	+0.034**	+0.013	+0.021	+0.022*	+0.028**
	(0.016)	(0.016)	(0.012)	(0.013)	(0.013)	(0.014)
Proportional, remote work held constant	+0.039**	+0.019	+0.004	+0.013	+0.013	
	(0.017)	(0.014)	(0.012)	(0.012)	(0.012)	
Proportional, common sample	+0.043***	+0.038**	+0.016	+0.018	+0.023*	+0.025*
	(0.016)	(0.016)	(0.013)	(0.014)	(0.014)	(0.014)
Placebo 2016-19, proportional	-0.027**	+0.000	-0.013	+0.011	+0.003	+0.006
	(0.013)	(0.015)	(0.012)	(0.015)	(0.013)	(0.011)

Notes: Points: percentage points per year. Proportional: change in the log unemployment rate per year (multiply by 100 for approximate percent). Major group effects, the placebo-window change and the COVID dip controlled. Remote-work potential is the Dingel and Neiman (2020) teleworkability share, entered as a control and, in its own column, as a predictor. Within-group shuffle tests (1,000 draws), two-sided p, points: Opportunity Data AI Exposure Index 0.005; Felten 0.86; Eloundou GPT-4 0.39; Eloundou human 0.80; BLS 0.65. Proportional: Opportunity Data AI Exposure Index <0.001; Felten 0.011.

4.4 Inside digital work

If capability is close to sufficient for digitally mediated tasks, capability ratings should stop separating digital occupations from one another. I split the occupations at the employment-weighted median of digital intensity. In the more digital half (186 groups), the four capability and usage measures retain 77 to 84 percent of their full-sample standard deviation, so they still rank these occupations differently; none of those rankings predicts the unemployment change at the 5 percent level (+0.03 to +0.08; p 0.06 to 0.34). The Opportunity Data AI Exposure Index predicts it at +0.21 ($p<0.001$), through the human contact (+0.18) and physical anchor (+0.17) dimensions (Figure 2). In the less digital half the capability measures predict the wrong sign: occupations they rate more exposed saw unemployment fall relative to their groups (-0.18 to -0.13; p 0.002 to 0.07).

Figure 2. Unemployment change per standard deviation, by half of the workforce split on digital intensity

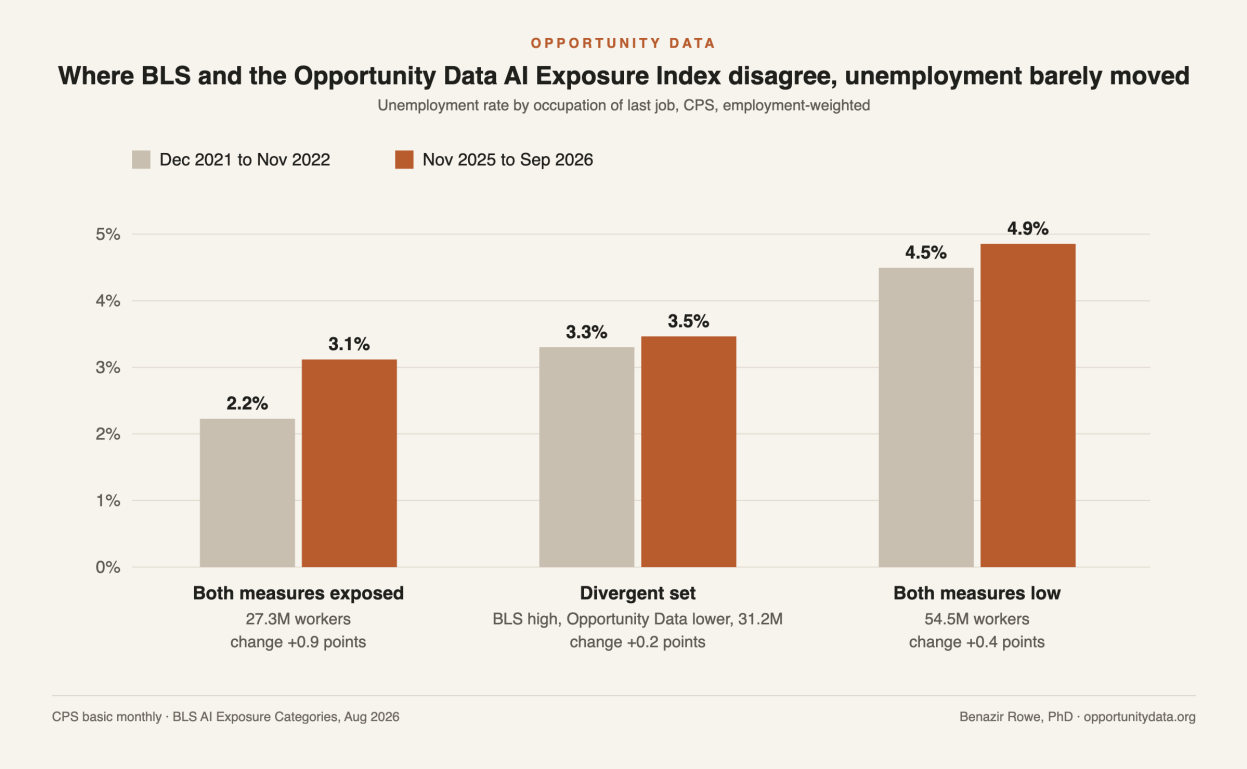


Notes: Estimates with 90 percent confidence intervals; major group effects, the placebo-window change and the COVID dip are controlled. Spread ratios are in Appendix D.

4.5 Where BLS and the Opportunity Data AI Exposure Index disagree

Rowe (2026) identified occupations that BLS rates High or Very high and the Opportunity Data AI Exposure Index rates two or more tiers lower, concentrated in teaching, nursing, counseling, front-line sales and service and office support. Usage-based inputs rate these occupations highly exposed because their workers use AI heavily; the structural measure rates them protected because the work is face to face. Figure 3 shows the outcome. Unemployment in this divergent set (31.2 million workers) rose 0.2 points, from 3.3 to 3.5 percent, against 0.9 points where both measures rate occupations exposed and 0.4 where both rate them low. OEWS employment in the divergent set grew 2.7 percent from 2022 to 2025, faster than its 1.2 percent from 2016 to 2019, while growth in the jointly exposed set slowed from 5.1 to 2.6 percent.

Figure 3. Unemployment by agreement between BLS and the Opportunity Data AI Exposure Index



Notes: Employment-weighted unemployment rate by occupation of last job, CPS. Retail salespersons and cashiers are a large share of the divergent set; cashier employment was declining before generative AI.

4.6 Outcomes no measure predicts

Young workers. Across occupations, every measure predicts slower employment growth at ages 22 to 27 than at 35 to 54 in the same occupation, by 1.63 to 1.79 points a year per standard deviation, with no such pattern in the placebo window. This reproduces in public survey data the direction of the payroll-data finding of Brynjolfsson, Chandar and Chen (2026). The pattern is between sectors (computer, business and office work against the rest): within major groups no measure is significant (all $p>0.33$), and the across-group estimates weaken at stricter sample floors (Appendix E).

Employment. Within major groups the Opportunity Data AI Exposure Index predicts slower OEWS employment growth (-0.56, $p=0.067$) until the knowledge-sector share is controlled (-0.20, $p=0.533$): most of that association reflects the 2022 to 2024 rate cycle and technology-sector retrenchment. The other measures range from -0.07 to +0.81.

Wages. Median wages grew faster in exposed occupations on every measure once trends are controlled (+0.24 to +0.52 points a year), the Opportunity Data AI Exposure Index the smallest at +0.24, consistent with rising earnings and stable unemployment in exposed occupations in US data (Chen et al. 2025) and with small earnings effects in Danish registers (Humlum and Vestergaard 2026).

4.7 Year by year

Rolling 12-month windows show that every measure predicted the 2024 rise in unemployment, the capability measures most strongly (peak estimates +0.31 to +0.47). Since September 2025, 1 of 40 year-over-year estimates across the eight measures are significant. The evidence does not yet show capability measures losing predictive power over time; it shows that their cumulative post-2022 association is near zero (Appendix F).

5. Robustness

Section 4.3 and Table 4 report the scale, remote-work, starting-level, common-sample and shuffle tests. The unemployment result also holds with the COVID dip controlled (+0.11, p=0.014), and the head-to-head result holds against each competitor. The employment result for the composite does not survive the sector control; the human contact dimension survives it (-0.59, p=0.041) but not an unweighted regression. Dropping computer and mathematical occupations, health occupations, or food and personal service occupations leaves the composite's employment estimate negative and insignificant in every case. Appendix E reports every variant.

6. Discussion

The evidence supports a narrow claim and rules out a broad one. The narrow claim: on the outcome where the five measures disagree, unemployment within occupational groups, the structure of work has predicted what happened since 2022 more strongly than 2023-vintage capability ratings, it is the only measure that survives a control for remote-work potential, and inside digital work the capability ratings no longer discriminate. That is what one would expect if capability ratings describe the model generation they were collected on and if, as capability approaches sufficiency for digital tasks, face-to-face contact and physical presence become what separates occupations. The divergent set gives the same answer from another direction: occupations that use AI heavily but are anchored by human contact have not seen the unemployment rise that usage data would predict. Exposure through use is not exposure to replacement.

The broad claim is not supported: no measure forecasts job loss. None predicts employment or wages once sector composition is controlled for, and the young-worker shift is visible across sectors but not within them. The unemployment result is one outcome from one survey, significant at conventional levels among roughly a dozen outcome windows; the shuffle test, the flat placebo and the remote-work control make chance and the most obvious confounder less likely explanations, but the forward test is the decisive check.

Limitations. Unemployment by occupation of last job reflects separations and slow re-hiring of experienced workers, not entry, where other studies locate the effect on young workers (Atkinson and Yamco 2026; Brynjolfsson, Chandar and Chen 2026). CPS occupation cells are small and occupation coding adds noise; robust standard errors do not account for correlation among related occupation groups. Controlling for the late-2022 unemployment rate weakens the composite (Section 4.3). The sector share is measured in 2025, inside the outcome window. OEWS is not designed as a time series. The Opportunity Data AI Exposure Index was constructed in 2026, after part of the outcome window

could be observed, so only the registered forward test is fully out of sample. The author built one of the measures; to limit that conflict every measure is run through identical code and every result, including those unfavorable to the Opportunity Data AI Exposure Index, is reported.

7. Forward test

The outcome windows end in September 2026, months after the April 2026 frontier on which the Opportunity Data AI Exposure Index was designed. Scores are frozen at the September 2026 release (sha256 7ce0a6a62dc3ceba...), and six hypotheses are registered to be scored on the CPS through September 2027 and September 2028 and on OEWS May 2026 and May 2027, with results published whatever they show. The registration and its frozen inputs were RFC 3161 timestamped on November 6, 2026 (freetsa.org, serial 0x09D16E2E) and are included in the replication package (Appendix G).

8. Conclusion

Of five AI exposure measures tested against US labor-market outcomes since late 2022, the one that measures the structure of work rather than the capability of AI is the strongest predictor of the rise in unemployment within occupational groups and the only one that survives a control for remote-work potential; the capability ratings no longer sort occupations inside digital work. No measure predicts employment or wages. Measures used to steer public money should be tested against outcomes on a schedule set in advance; this paper sets one.

Data and code availability. All data are public. The replication package (code, harmonized inputs, monthly CPS cells, every result table, checksums) reproduces every number in this paper in about one minute; the Felten et al. and Dingel and Neiman files are downloaded from their sources with checksum checks. https://opportunitydata.org/ai-exposure-outcome-validation.html

Disclosure. Opportunity Data is an independent research platform. The author built the Opportunity Data AI Exposure Index, which is one of the five measures tested; to limit that conflict, every measure is run through identical code, every outcome and specification is published including those where the Opportunity Data AI Exposure Index does not outperform, and the forward test is fixed in advance.

References

Atkinson, T. and S. Yamco (2026). Young workers' employment drops in occupations with high AI exposure. Federal Reserve Bank of Dallas, January 6, 2026.

Brynjolfsson, E., B. Chandar and R. Chen (2026). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. Stanford Digital Economy Lab, revised August 2026 (first version August 2025).

Bureau of Labor Statistics (2026). AI exposure categories, 2025-35 employment projections. August 27, 2026. https://www.bls.gov/emp/publications/ai-exposure-categories.htm

Chen, Kane, Kozlowski, Kunievsky and Evans (2025). The (short-term) effects of large language models on unemployment and earnings. arXiv:2509.15510.

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

Eckhardt, S. and N. Goldschlag (2025). All quiet on the AI front. Economic Innovation Group, August 11, 2025.

Eloundou, T., S. Manning, P. Mishkin and D. Rock (2024). GPTs are GPTs: Labor market impact potential of LLMs. Science 384(6702): 1306-1308. Data: github.com/openai/GPTs-are-GPTs.

Fairlie, R. W. and J. Wu (2026). The early impacts of AI on employment among recent college graduates. NBER Working Paper 35796.

Felten, E., M. Raj and R. Seamans (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. Strategic Management Journal 42(12): 2195-2217.

Felten, E., M. Raj and R. Seamans (2023). How will language modelers like ChatGPT affect occupations and industries? arXiv:2303.01157. Data: github.com/AIOE-Data/AIOE.

Frank, M. R., Y.-Y. Ahn and E. Moro (2023). AI exposure predicts unemployment risk. arXiv:2308.02624.

Frank, M. R., A. Javadian Sabet, L. Simon, S. H. Bana and R. Yu (2026). AI-exposed jobs deteriorated before ChatGPT. arXiv:2601.02554.

Gimbel, M., M. Kinder, J. Kendall and M. Lee (2025). Evaluating the impact of AI on the labor market: Current state of affairs. The Budget Lab at Yale, October 2025.

Handa, K., et al. (2025). Which economic tasks are performed with AI? Evidence from millions of Claude conversations. arXiv:2503.04761.

Humlum, A. and E. Vestergaard (2026). Still waters, rapid currents: Early labor market transformation under generative AI. NBER Working Paper 33777, revised March 2026.

Lund, Euyang, Munyikwa and Fadaee (2026). AI exposure scores: What they measure, what they miss, and what comes next. arXiv:2606.23633.

Massenkoff, M. and P. McCrory (2026). Labor market impacts of AI: A new measure and early evidence. Anthropic, March 5, 2026.

Rowe, B. (2026). The Opportunity Data AI Exposure Index: Occupational AI exposure after the frontier break of 2026. Opportunity Data working paper, September 2026.

Tomlinson, K., S. Jaffe, W. Wang, S. Counts and S. Suri (2025). Working with AI: Measuring the applicability of generative AI to occupations. arXiv:2507.07935.

Tucker, L. C. (2026). You're (not) hired: Artificial intelligence and early career hiring in the Quarterly Workforce Indicators. Center for Economic Studies Working Paper 26-27, U.S. Census Bureau.

U.S. Census Bureau. Current Population Survey basic monthly public-use files, January 2015 to September 2026.

U.S. Bureau of Labor Statistics. Occupational Employment and Wage Statistics, national estimates, May 2016 to May 2025.

Appendix A. Data sources and vintages

Table A1. Inputs

Input	Vintage	Use
CPS basic monthly public-use files	January 2015 to September 2026 (October 2025 not collected)	Employment, unemployment by occupation of last job, age; composite weight (PWCMPWGT); occupation PEIO1OCD/PTIO1OCD, columns 860-863
OEWS national, detailed occupations	May 2016, 2017, 2018, 2019, 2022, 2023, 2024, 2025	Employment (TOT_EMP), median annual wage (A_MEDIAN); 2016-2024 files are Internet Archive copies of the BLS special-requests archives
OEWS all-data file	May 2025	National occupation employment by NAICS sector, for the knowledge-sector share
Opportunity Data AI Exposure Index	September 2026 release	Composite and three dimensions, SOC 2018
Felten, Raj and Seamans LM-AIOE	AIOE-Data GitHub release	SOC 2010; mapped to SOC 2018
Eloundou et al. occ_level.csv	openai/GPTs-are-GPTs	dv_rating_beta (GPT-4) and human_rating_beta, averaged across O*NET-SOC detail codes
BLS AI Exposure Categories	August 27, 2026	Low, Moderate, High, Very high coded 1 to 4
Crosswalks	BLS SOC 2010-2018 (2017); Census occupation codes 2010-2018 (September 2019)	Occupation harmonization

Appendix B. Occupation harmonization

Occupation codes are treated as nodes of a graph whose edges are the official crosswalk links (SOC 2010 to SOC 2018 for OEWS; Census 2010 to Census 2018 occupation codes for CPS). Each connected component is one occupation group, so any code that split or merged is pooled with all codes it touched, and every window compares the same groups. OEWS hybrid codes used in 2019 that appear in neither classification join the group of their broad (then minor) occupation. A group-year is set to missing if any member occupation has suppressed employment in that year. Measure scores are attached at the SOC 2018 detailed level and averaged within each group with employment weights; Felten et al. scores are mapped from SOC 2010 through the crosswalk, with broad-group and then minor-group means for residual codes the source omits; Census occupation codes that cover several SOC codes take the employment-weighted mean of those codes. The result is 738 OEWS groups with employment in May 2022 and May 2025 and 525 CPS groups.

Appendix C. Outcome construction and sample

CPS windows are 12-month averages of monthly weighted counts. Unemployment by occupation counts experienced unemployed persons (labor force status 3 or 4 with a reported last occupation) against employed persons in the same group. The young-worker gap is kept where at least 30 unweighted respondents aged 22 to 27 appear in both windows (246 groups). Changes are divided by the years between window midpoints (3.9 years post-2022, 3.0 years placebo). OEWS top-coded medians are treated as missing. The knowledge-sector share is employment in NAICS 51, 52, 54 and 55 divided by total employment in the occupation, May 2025. The COVID dip is log employment change

from May 2019 to May 2022 (OEWS) or from the placebo base window to the post-2022 base window (CPS). Divergent-set tiers use the published Opportunity Data AI Exposure Index cutoffs (0.45, 0.60, 0.70); a group is divergent if its BLS category is High or Very high and its tier is two or more lower.

Appendix D. Every outcome and specification

Table D1. Coefficient per standard deviation, all outcomes, measures and specifications

Outcome	Specification	Composite	Digital	Contact	Physical	Felten	Elound. GPT-4	Elound. human	BLS
cps emp post	bivariate	+0.23	+0.17	-0.05	+0.33	+0.31	+0.14	+0.07	+0.22
cps emp post	fe plus controls	-0.51	-0.65*	-0.54	+0.17	+0.16	-0.24	-0.31	+0.36
cps emp post	major group fe	-0.57	-0.70	-0.09	-0.18	-0.32	-0.61	-0.84*	-0.04
cps emp pre	bivariate	+0.35	+0.65***	-0.26	+0.24	+0.41	+0.30	+0.35	+0.11
cps emp pre	major group fe	+0.54	+0.67*	+0.06	+0.30	+0.37	+0.02	+0.43	-0.08
cps urate post	bivariate	+0.06***	+0.01	+0.08***	+0.04*	+0.02	+0.02	+0.02	+0.02
cps urate post	fe plus controls	+0.11**	+0.03	+0.12***	+0.10*	+0.01	-0.03	-0.01	-0.02
cps urate post	major group fe	+0.10**	-0.01	+0.11***	+0.09**	+0.01	-0.03	-0.01	-0.01
cps urate pre	bivariate	+0.08***	+0.09***	-0.07***	+0.11***	+0.17***	+0.14***	+0.15***	+0.16***
cps urate pre	major group fe	-0.01	+0.02	-0.04	-0.02	+0.09*	+0.09**	+0.10**	+0.10**
cps young gap post	bivariate	-1.64***	-1.26**	-0.51	-1.64***	-1.79***	-1.75***	-1.76***	-1.63***
cps young gap post	fe plus controls	-0.42	-0.40	+0.08	-0.52	-0.99	-0.85	-0.89	-0.55
cps young gap post	major group fe	-0.16	-0.01	+0.05	-0.39	-0.68	-0.86	-0.53	-0.25
cps young gap pre	bivariate	+0.96*	+0.44	+0.79	+0.87	-0.04	+0.26	-0.20	-0.00
cps young gap pre	major group fe	+1.76*	+0.20	+1.51**	+2.02*	+0.87	+0.50	-0.44	+0.31
oews emp 2224	bivariate	-0.04	+0.27	-0.44	+0.03	+0.22	-0.45	-0.25	-0.21
oews emp 2224	fe plus controls	-0.52*	-0.23	-0.67**	+0.13	+0.60	-0.29	+0.19	+0.32
oews emp 2224	major group fe	-0.29	+0.00	-0.64**	+0.19	+0.39	-0.33	+0.17	+0.30
oews emp 2325	bivariate	-0.30	+0.20	-0.87***	-0.06	+0.18	-0.50*	-0.31	-0.29
oews emp 2325	fe plus controls	-0.62*	-0.13	-0.87***	+0.10	+0.35	-0.28	-0.06	+0.18
oews emp 2325	major group fe	-0.45	-0.00	-0.85**	+0.12	+0.22	-0.33	-0.08	+0.15
oews emp post	bivariate	-0.15	+0.22	-0.64*	+0.03	+0.23	-0.45	-0.23	-0.22
oews emp post	fe plus controls	-0.56*	-0.22	-0.75***	+0.14	+0.54	-0.29	+0.09	+0.23
oews emp post	major group fe	-0.37	-0.07	-0.72**	+0.18	+0.37	-0.35	+0.05	+0.20
oews emp pre	bivariate	+0.17	+0.35	+0.16	-0.11	-0.10	-0.38	-0.26	-0.32
oews emp pre	fe plus controls	-0.32	-0.05	-0.02	-0.58	-0.71	-0.81*	-0.17	+0.00
oews emp pre	major group fe	-0.10	+0.13	-0.02	-0.36	-0.36	-0.57	-0.05	+0.09
oews wage post	bivariate	-0.30***	-0.30**	-0.17	-0.18	-0.17	-0.16	-0.12	-0.16
oews wage post	fe plus controls	+0.24	+0.18	-0.08	+0.44**	+0.42*	+0.36*	+0.52***	+0.37*
oews wage post	major group fe	-0.03	-0.08	-0.11	+0.17	+0.05	+0.09	+0.25	+0.11
oews wage pre	bivariate	-0.64***	-0.72***	-0.04	-0.52***	-0.66***	-0.71***	-0.59***	-0.57***
oews wage pre	fe plus controls	+0.05	+0.00	+0.06	+0.02	-0.07	-0.24	-0.01	+0.02
oews wage pre	major group fe	-0.38*	-0.44*	+0.02	-0.37*	-0.65**	-0.66***	-0.48***	-0.42*

Notes: Composite, Digital, Contact and Physical are the Opportunity Data AI Exposure Index and its dimensions. Outcomes ending "pre" are placebo windows. *** p<0.01, ** p<0.05, * p<0.1. Standard errors are in the replication package (regression_results.csv).

Table D2. Unemployment, 2022-26, by half of the workforce split on digital intensity

Measure	More digital half	Less digital half	Spread ratio, more digital half
Opportunity Data AI Exposure Index	+0.21*** (0.04)	-0.05 (0.06)	0.91
Human contact dimension	+0.18*** (0.05)	+0.07 (0.06)	1.05
Physical anchor dimension	+0.17*** (0.03)	-0.06 (0.07)	0.89
Digital intensity dimension	+0.09*** (0.03)	-0.08* (0.05)	0.48
Felten, Raj and Seamans LM-AIOE	+0.05 (0.05)	-0.13* (0.07)	0.77
Eloundou et al., GPT-4 ratings	+0.08* (0.04)	-0.14** (0.06)	0.84
Eloundou et al., human ratings	+0.05 (0.04)	-0.18*** (0.07)	0.84
BLS AI Exposure Categories	+0.03 (0.03)	-0.18*** (0.06)	0.84

Notes: Spread ratio: standard deviation of the measure inside the half divided by its standard deviation in the full sample.

Table D3. Agreement groups

Group	Workers	Unemployment base	Latest	2016	2019	OEWS growth 2022-25	2016-19
both exposed	27.3M	2.23	3.12	3.02	2.37	2.6%	5.1%
divergent	31.2M	3.30	3.46	4.01	3.47	2.7%	1.2%
both low	54.5M	4.49	4.85	5.88	4.42	4.4%	4.1%
other	45.0M	2.64	3.16	3.19	2.38	9.8%	8.0%

Appendix E. Robustness

Table E1. Robustness variants

Outcome	Variant	Composite	Felten et al. LM-AIOE	Eloundou et al. GPT-4	Eloundou et al. human	BLS categories
cps emp post	add covid recovery	-0.50 (0.44)	+0.13 (0.49)	-0.24 (0.41)	-0.35 (0.48)	+0.34 (0.37)
cps emp post	main	-0.51 (0.43)	+0.16 (0.48)	-0.24 (0.40)	-0.31 (0.47)	+0.36 (0.36)
cps urate post	add covid recovery	+0.11** (0.04)	+0.01 (0.05)	-0.03 (0.04)	-0.01 (0.04)	-0.02 (0.04)
cps urate post	main	+0.11** (0.04)	+0.01 (0.05)	-0.03 (0.04)	-0.01 (0.04)	-0.02 (0.04)
cps young gap post	min n 30 bivariate	-1.64*** (0.48)	-1.79*** (0.53)	-1.75*** (0.41)	-1.76*** (0.50)	-1.63*** (0.45)
cps young gap post	min n 30 major group fe	-0.16 (1.04)	-0.68 (1.05)	-0.86 (0.89)	-0.53 (1.07)	-0.25 (0.90)
cps young gap post	min n 50 bivariate	-1.42*** (0.49)	-1.54*** (0.55)	-1.57*** (0.42)	-1.55*** (0.51)	-1.38*** (0.47)
cps young gap post	min n 50 major group fe	+0.24 (1.08)	+0.08 (1.08)	-0.40 (0.91)	+0.23 (1.08)	+0.62 (0.92)
cps young gap post	min n 80 bivariate	-1.31** (0.53)	-1.29** (0.57)	-1.45*** (0.45)	-1.35** (0.54)	-1.22** (0.50)
cps young gap post	min n 80 major group fe	+0.79 (1.20)	+1.42 (1.03)	+0.38 (0.99)	+1.35 (1.13)	+1.70* (0.92)
oews emp post	add covid and sector	-0.20 (0.32)	+0.81* (0.49)	-0.07 (0.34)	+0.29 (0.43)	+0.43 (0.33)
oews emp post	add covid recovery	-0.57* (0.31)	+0.51 (0.50)	-0.29 (0.33)	+0.08 (0.44)	+0.22 (0.34)
oews emp post	drop computer math 15	-0.18 (0.32)	+0.83* (0.49)	-0.05 (0.34)	+0.30 (0.43)	+0.44 (0.33)
oews emp post	drop food personal 35 39	-0.10 (0.33)	+0.96* (0.52)	-0.06 (0.35)	+0.22 (0.44)	+0.53 (0.34)
oews emp post	drop health 29 31	-0.30 (0.35)	+0.71 (0.57)	-0.04 (0.39)	-0.01 (0.48)	+0.16 (0.40)
oews emp post	main	-0.56* (0.31)	+0.54 (0.48)	-0.29 (0.32)	+0.09 (0.43)	+0.23 (0.33)
oews emp post	unweighted	-0.68* (0.39)	+0.36 (0.55)	-0.47 (0.41)	+0.32 (0.36)	+0.26 (0.38)
oews emp pre placebo	add covid and sector	-0.33 (0.43)	-0.71 (0.52)	-0.80* (0.46)	-0.17 (0.59)	+0.00 (0.53)

Notes: Composite is the Opportunity Data AI Exposure Index. Same units and notation as Table 2.

Appendix F. Year by year

Table F1. Year-over-year change in unemployment rate per standard deviation, rolling 12-month windows

Window ending	Composite	Felten et al. LM-AIOE	Eloundou et al. GPT-4	Eloundou et al. human	BLS categories
2022-12	+0.33*	+0.15	+0.36**	+0.20	+0.09
2023-03	+0.27*	+0.03	+0.22*	+0.03	+0.03
2023-06	+0.05	-0.11	+0.06	-0.23*	-0.10
2023-09	-0.01	-0.15	-0.12	-0.25**	-0.11
2023-12	+0.10	-0.09	-0.03	-0.12	-0.06
2024-03	+0.19*	+0.21	+0.10	+0.12	+0.12
2024-06	+0.24**	+0.29**	+0.17*	+0.25**	+0.21**
2024-09	+0.34***	+0.47***	+0.25**	+0.39***	+0.31***
2024-12	+0.22**	+0.43***	+0.17	+0.29***	+0.31***
2025-03	+0.15	+0.09	-0.03	+0.13	+0.14
2025-06	+0.06	-0.08	-0.14	+0.00	+0.00
2025-09	-0.06	-0.23	-0.15	-0.14	-0.12
2025-12	-0.05	-0.31**	-0.17	-0.19	-0.17
2026-03	-0.08	-0.12	-0.12	-0.15	-0.14
2026-06	+0.07	+0.05	+0.02	+0.00	-0.03
2026-09	+0.14	-0.02	-0.06	-0.04	-0.10

Notes: Each row compares the 12 months ending in the stated month with the 12 months before, within major groups. Composite is the Opportunity Data AI Exposure Index.

Appendix G. Registered forward test

Frozen scores: ai-exposure-soc-data.js, sha256 7ce0a6a62dc3ceba22e2d69d276d85da580bd7629177b3c2197f8cd81db8fad5. Windows: CPS base November 2025 to September 2026, outcomes October 2026 to September 2027 and October 2027 to September 2028; OEWS base May 2025, outcomes May 2026 and May 2027. One-sided tests at the 5 percent level within occupational major groups, employment-weighted, robust standard errors. Every hypothesis is reported at each scoring date whatever the result, alongside the same statistics for each competing measure.

H1. Unemployment by occupation of last job rises faster in occupations with higher Opportunity Data AI Exposure Index scores, within occupational groups.

H2. The H1 result holds with each competing measure in the same regression.

H3. Unemployment in the divergent set rises less than in the jointly exposed set.

H4. OEWS employment growth from May 2025 is slower in higher-scoring occupations, with sector controls.

H5. Inside the more digital half of the workforce, the H1 result holds for the Opportunity Data AI Exposure Index and the four capability and usage measures add no predictive power.

H6 (exploratory). Within occupational groups, employment of 22 to 27 year-olds falls relative to 35 to 54 year-olds in higher-scoring occupations.