The Recent Evolution of AI-Related Labor Demand*

Kevin Rinz

Federal Reserve Bank of Cleveland

August 19, 2026

Abstract

Mentions of AI in US job ads have increased substantially since the beginning of 2024. Mentions have come in occupations more intensely exposed to the capabilities of large language models. Over the same period, the quantity of job postings has stabilized and posted wages have risen. Both hires and separations have increased in exposed occupations. Relaized wages for new hires have risen comparably to posted wages, while wages of continuously employed workers have declined, possibly because exposed occupations have seen labor markets become less tight.

Keywords: artificial intelligence, labor demand, skills, job postings
JEL codes: J23, J24, O33.

*I thank Kevin Jacobson and Carol Moseley for helpful research assistance. The analysis underlying this paper was conducted with the help of artificial intelligence, specifically Claude Opus 4.8, Claude Opus 5, and Claude Fable 5. Contact: kevin.rinz@clev.frb.org

Artificial intelligence (AI) capabilities are expanding rapidly, raising questions about potential labor market consequences. Some research on this topic suggests that demand for labor has largely held up in aggregate (Alekseeva et al., 2021; Acemoglu et al., 2022; Hampole et al., 2025; Humlum and Vestergaard, 2025), though much of this work mainly considers a period when frontier models were meaningfully less capable than they are today.[1] More recent work points to possible exceptions for narrower groups, like early career workers (Brynjolfsson et al., 2025; Tucker, 2026) or freelancers (Hui et al., 2024). In this paper, I examine the recent evolution of AI-related labor demand, considering how AI capabilities relate to the content and quantity of job ads, posted and realized wages, and employment flows.

After gradually trending up for years, mentions of AI-related terms in online ads for full-time jobs collected by Lightcast spiked in 2024Q3 (Figure 1a). This spike consisted almost entirely of mentions of the term "artificial intelligence" (shown in green) and occurred mostly in ads for jobs in computer and mathematical occupations (light green) like software developers, data scientists, and computer and information research scientists.[2] These mentions quickly receded. In aggregate, they were replaced by mentions of the acronym "AI" in combination with terms that are more generic or suggestive of company or product branding rather than tied to specific skills (shown in orange).[3] Aggregate mentions of specific AI skills have evolved more continuously over this period and have accelerated in recent months. Recent months have also seen increased mentions of AI in the context of disclosing employer use of AI in hiring processes (shown in purple). See Appendix B for discussion of how these different types of mentions are identified. That appendix also discusses analysis comparing the content of AI and non-AI ads over time, which shows modest shifts alongside the spike

[1] For example, Claude Code and Codex became widely available in May 2025. Claude Fable and GPT-5.6 Sol, each of which have had access restrictions imposed by the federal government, were released in June and July 2026, respectively.

[2] Figure D5 shows top occupations mentioning AI-related terms over the last year and their rates of mentioning AI during the 2024Q3 spike.

[3] Examples include terms like AI company, AI-native, predictive AI, harness AI, AI products, and similar constructions.

in AI mentions that also revert.

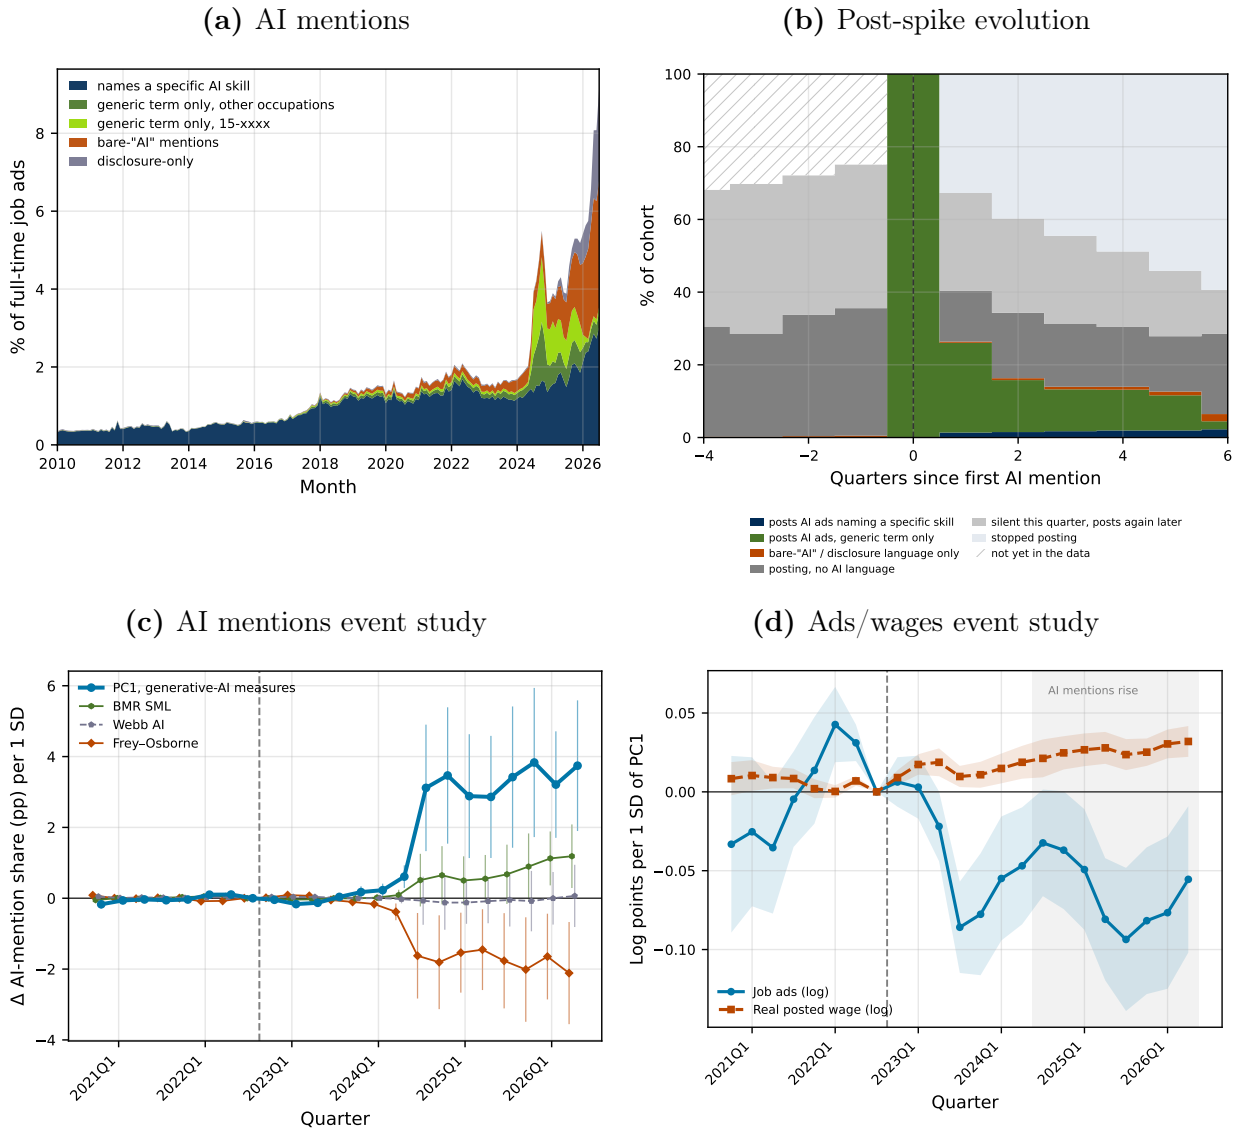
Only 40 percent of the particular employer-occupations that newly mentioned "artificial intelligence" during 2024Q3 posted another ad the next quarter; of those, about 35 percent no longer mentioned that term or any other related to AI (Figure 1b). The shares posting ads and mentioning "artificial intelligence" continued to decline in subsequent quarters, while only about 2 percent of employer-occupations that initially mentioned "artificial intelligence" during 2024Q3 mention a more specific AI-related term in a job ad six quarters later. The most common occupations making these mentions and their subsequent reversion together suggest more that AI skills are rapidly becoming table stakes than that employers were experimenting briefly and decided AI was not important.

The short-lived nature of the spike in "artificial intelligence" mentions and its concentration within tech occupations suggests that the spike represented a relatively superficial relabeling of jobs. The reversion of the language used even within employer-occupation cells and the fact that survey-measured AI adoption has grown at a faster pace than AI mentions in job ads since the spike (Figure D6) points to potential learning by employers that AI skills are becoming table stakes, and it may not be necessary to discuss them explicitly in ads.

Increased mentions of AI-related terms have tended to come in ads for jobs that are more exposed to AI by certain measures, specifically measures related to large language models (LLMs). Recent work has quantified occupational exposure to AI based on the ability of LLMs to perform relevant tasks (Felten et al., 2021, 2023; Eloundou et al., 2024; Handa et al., 2025; Eisfeldt et al., 2025, I use the first principal component of these measures), uses described in AI patents (Webb, 2020), and the probability pre-LLM machine learning tools could fully automate jobs (Frey and Osborne, 2017) or automate constituent tasks (Brynjolfsson et al., 2018).[4] I incorporate these exposure measures into event study regressions of

[4]See Appendix A for a more detailed description of each measure (Table A1) and pairwise correlations between them and comparisons of most exposed occupations (Table A2). Figure A1 reproduces the event study estimates from Figure 1 using each individual exposure measure.

Figure 1: AI Exposure and Job Ads



Notes: Lightcast online job ads for full-time positions. Panel (a) plots the monthly share of ads containing each type of AI reference: a skill drawn from a curated list of AI-related terms; the phrase "artificial intelligence" as the only such term, shown separately for computer and mathematical occupations; generic uses of the acronym "AI"; and language disclosing an employer's use of AI in hiring. The categories are mutually exclusive, so the bands sum to the overall mention rate. Panel (b) follows the employer-occupation pairs whose first AI mention was the generic phrase in 2024Q3 and classifies what each posted in the quarters around it. Panels (c) and (d) report β_t from equation 1 with 95 percent confidence intervals, estimated on the 628 occupations (636 for posted wages) that clear the size screen and are scored by every component measure, with 2022Q3 omitted and exposure measured in standard deviations; standard errors are clustered on occupation. Panel (c) compares the first principal component of the exposure measures based on large language models with the three measures excluded from it.

the form

$$y_{ot} = \sum_{t \neq 2022q3} \beta_t \left(\text{Exposure}_o \times 1[\text{Quarter=t}]\right) + \alpha_o + \delta_t + \epsilon_{ot} \quad (1)$$

where outcomes y_{ot} include the share of job ads mentioning AI-related terms, the total number of job ads posted, and posted wages, $Exposure_o$ is a demeaned exposure measure in standard deviation units, α_o represents occupation fixed effects and δ_t represents quarter fixed effects.[5] The coefficients β_t represent the change in outcome y_{ot} from the quarter before ChatGPT's release (2022Q3) to quarter t for each standard deviation of AI exposure in occupation o.[6] Figure 1c shows that the share of job ads mentioning AI-related terms increased by about 3 percentage points per standard deviation of occupational exposure to LLM capabilities in 2024Q3, increasing to nearly 4 percentage points in subsequent quarters, suggesting that LLM exposure has some relevance in the current labor market. Other exposure measures have more muted or even negative correlations with AI mentions in job ads.

Alongside increased mentions of AI, the relative quantity of ads in more LLM-exposed occupations roughly stabilizes, following a period of decline that began prior to ChatGPT's release (Figure 1d). Over the same period, real posted wages continue rising on a trajectory that began right after ChatGPT's release.[7] The increases in AI mentions and posted wages, in particular, appear to be driven by the highest-exposure occupations (Figure D2); excluding tech occupations from the analysis does not meaningfully alter the pattern of estimates

[5] I use a balanced panel of occupations, and standard errors are clustered at the occupation level. For the posted wage analysis, I include 652 occupations with at least 25 ads that contain wage information every quarter. For AI mentions and job postings analysis, I include 643 occupations with at least 125 job ads every quarter. Estimates using the first principal component are identified from the 636 and 628 of those occupations, respectively, that are scored by all five of its component measures.

[6] This formulation is intended to peg estimates to a baseline immediately before a highly salient and relevant event, not to represent the causal effects of ChatGPT on these outcomes. Inclusion of additional controls for education and experience requirements makes very little difference to the estimates presented in this paper. See Figure D1. Since education and experience requirements are potentially endogenous to the degree of AI exposure in an occupation, I omit these controls from my baseline analysis.

[7] The share of job ads that include wage information increases substantially over the period considered here, raising potential concerns about changing selection into wage posting. My wage estimates survive alternative specifications using alternative samples, reweighting, and bounding techniques meant to address different aspect of potential selection. See Appendix C for a detailed discussion.

reported here. While there is an ad-level AI posted wage premium (Figure D4), the wage increase in more LLM-exposed occupations is, if anything, clearer in ads not mentioning AI (Figure D3).

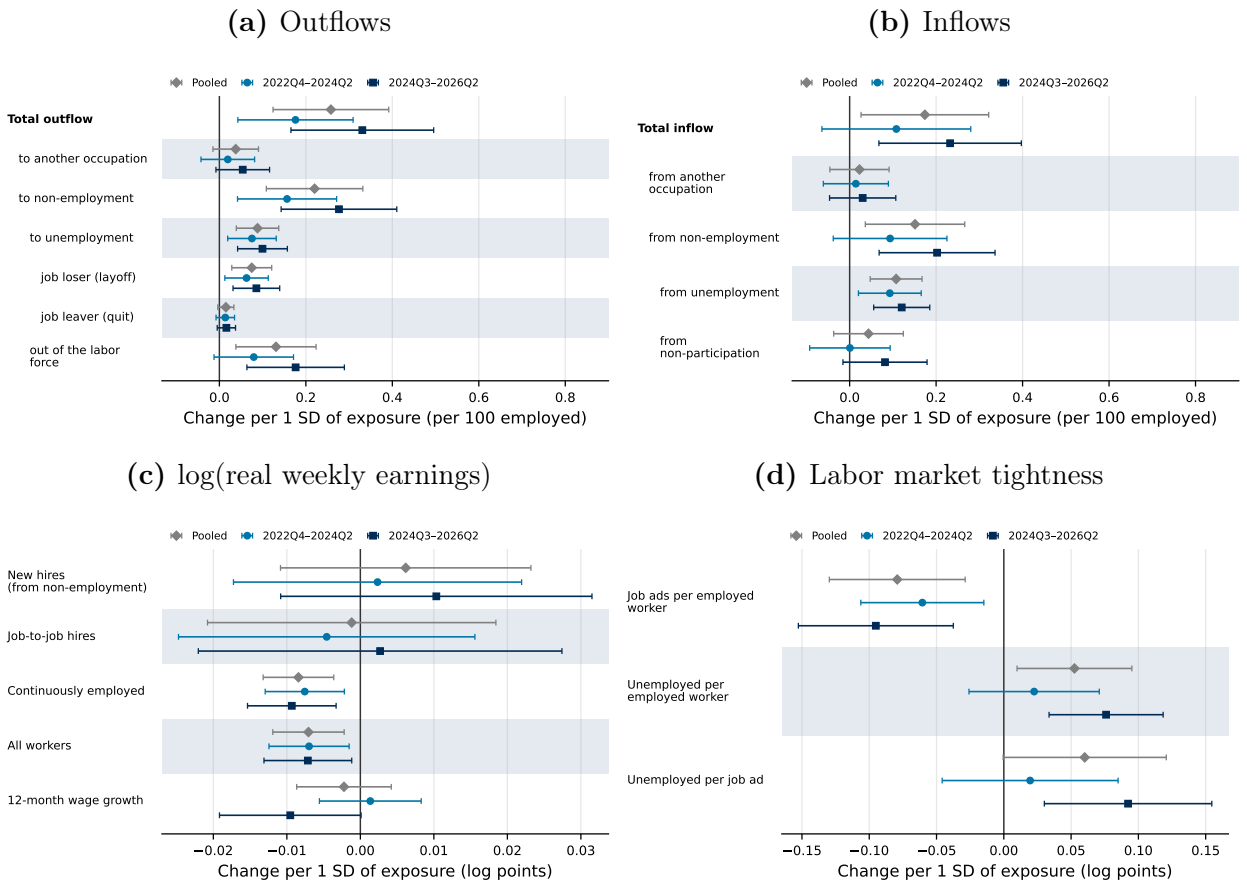
Job ads show increased discussion of AI alongside stable quantities of ads and rising posted wages. Together, these do not suggest collapsing labor demand, consistent with prior work. But posting ads is not the same as hiring, wage posting can be strategic and does not necessarily translate cleanly into realized wages, and tracking ads does not account for potential changes in separations or employers' desire to replace departed employees. These are all important considerations for understanding potential shifts in labor demand, and I turn to the Current Population Survey (CPS) to explore them further.

More LLM-exposed occupations saw increases in both outflows (Figure 2a) and inflows (Figure 2b) following the introduction of ChatGPT.[8] Outflows have been largely to non-employment, with most to non-participation. Among transitions to unemployment, almost all have been due to layoffs; there is no evidence of increased quits to unemployment, and a small but not statistically significant increase in job-to-job changes in occupation. Inflows are also primarily from non-employment, though most are from unemployment rather than non-participation. Point estimates for outflows are slightly larger, suggesting net flows out of more exposed occupations, but the difference is not statistically significant. The churn/gross flow rate is nearly 0.4 percentage points higher after the introduction of ChatGPT and about 0.5 percentage points higher during the period in which AI mentions in job ads rise.

Estimated changes in the realized real wages of new hires are positive, reaching about 1 percent since AI mentions spiked, but imprecise; the confidence intervals for changes in

[8]Due to smaller sample sizes, I use fewer, longer time periods in the CPS event study analysis: 2020Q4-2022Q3 (pre-ChatGPT), 2022Q4-2024Q2 (post-ChatGPT, pre-AI mention spike), and 2024Q3-2026Q2 (since AI mentions spiked). Regressions with labor flows as outcomes are estimated at the occupation-quarter level by weighted least squares, with flows expressed per 100 employed and cells weighted by pre-period employment. Ratios of unemployed workers to job ads or to employment are instead estimated by Poisson regression of the form $E\left[y_{ot} \mid \cdot\right] = \exp\left(\sum_p \beta_p \left(\text{Exposure}_o \times 1[t \in p]\right) + \alpha_o + \delta_t + \ln D_{ot}\right)$, where y_{ot} counts unemployed workers or job ads and the denominator D_{ot} enters as an offset, so β_p is the change in the log ratio per standard deviation of exposure. Poisson is preferable to least squares on the log ratio because 28 percent of occupation-quarter cells contain no unemployed respondents, and taking logs would drop exactly those cells.

Figure 2: AI Exposure, Employment Flows, and Earnings



Notes: Current Population Survey, 2020Q4–2026Q2. Points are difference-in-differences coefficients per standard deviation of AI exposure, measured relative to the pre-ChatGPT period (2020Q4–2022Q3), with 95 percent confidence intervals. Exposure is the first principal component of the exposure measures based on large language models, averaged to the Census occupation. The pooled estimate and the two post-period estimates come from separate regressions, so the pooled coefficient is approximately a duration-weighted average of the other two rather than a third estimate. Panels (a) and (b) use occupation-quarter cells built from month-to-month person links, expressed per 100 employed and weighted by pre-period employment; panel (c) uses individual weekly earnings records, in logs and deflated by the CPI-U; panel (d) combines unemployment counts from the CPS with job ads from Lightcast and is estimated by Poisson regression with the denominator as an offset. All specifications include occupation and quarter fixed effects and controls for age group, sex, race and ethnicity, and education. Standard errors are clustered on occupation.

realized wages generally contain the point estimates for changes in posted wages (Figure 2c).[9] Wage changes for workers making job-to-job transitions are closer to zero. Wages of continuously employed workers not changing jobs fall slightly. This may be related to the fact that more LLM-exposed occupations see tightness decline over the same period (Figure 2d), suggesting worse outside options/reduced upward pressure on wages for continuously employed workers.[10]

In the context of changes in employment flows, increased AI prevalence seems more clearly not to have dramatically reduced labor demand but may have weakened it modestly, with more LLM-exposed occupations having seen their labor markets become less tight recently, with unemployment rising relative to job postings. Exposure itself does seem to mean something in the current labor market, with mentions of AI more common in more exposed occupations, but going forward, directly tracking adoption by workers and firms may be more valuable than measuring the content of job ads. In occupations like data scientists, software developers, and web developers, the typical quarter over the last year has seen fewer mentions of AI in job ads than did 2024Q3. It seems unlikely that AI has become less important within these occupations since then, and measurement would ideally reflect that.

[9] Wage regressions are estimated at the individual level using a specification analogous to Equation 1 with controls for age, race/ethnicity, sex, and education added. Wages enter in logs and are adjusted for inflation using the CPI-U.

[10] Here, analysis of tightness uses new job ads posted during a period, which is more easily measured in Lightcast data at the occupation level, rather than the stock of current job openings, which is measured by the Job Openings and Labor Turnover Survey.

References

Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo (2022). "Artificial intelligence and jobs: Evidence from online vacancies." *Journal of Labor Economics*, 40(S1), pp. S293–S340.

Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska (2021). "The demand for AI skills in the labor market." *Labour Economics*, 71, p. 102,002.

Bick, Alexander, Adam Blandin, and David J. Deming (2024). "The rapid adoption of generative AI." NBER Working Paper 32966, National Bureau of Economic Research.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (2025). "Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence." Working paper, Stanford Digital Economy Lab.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2018). "What can machines learn, and what does it mean for occupations and the economy?" *AEA Papers and Proceedings*, 108, pp. 43–47.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019). "The effect of minimum wages on low-wage jobs." *Quarterly Journal of Economics*, 134(3), pp. 1405–1454.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux (1996). "Labor market institutions and the distribution of wages, 1973–1992: A semiparametric approach." *Econometrica*, 64(5), pp. 1001–1044.

Dube, Arindrajit, Daniele Girardi, Òscar Jordà, and Alan M. Taylor (2025). "A local projections approach to difference-in-differences." *Journal of Applied Econometrics*, 40(7), pp. 741–758.

Eisfeldt, Andrea L., Gregor Schubert, Bledi Taska, and Miao Ben Zhang (2025). "Generative AI and firm values." NBER Working Paper 31222, National Bureau of Economic Research.

Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock (2024). "GPTs are GPTs: Labor market impact potential of large language models." *Science*, 384(6702), pp. 1306–1308. ArXiv:2303.10130.

Felten, Edward W., Manav Raj, and Robert Seamans (2021). "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." *Strategic Management Journal*, 42(12), pp. 2195–2217.

Felten, Edward W., Manav Raj, and Robert Seamans (2023). "How will language modelers like ChatGPT affect occupations and industries?" Working paper. ArXiv:2303.01157.

Frey, Carl Benedikt and Michael A. Osborne (2017). "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change*, 114, pp. 254–280.

Hainmueller, Jens (2012). "Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies." *Political Analysis*, 20(1), pp. 25–46.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller (2025). "Artificial intelligence and the labor market." NBER Working Paper 33509, National Bureau of Economic Research.

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, et al. (2025). "Which economic tasks are performed with AI? evidence from millions of Claude conversations." Anthropic economic index, Anthropic. ArXiv:2503.04761.

Heckman, James J. (1979). "Sample selection bias as a specification error." *Econometrica*, 47(1), pp. 153–161.

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(6).

Humlum, Anders and Emilie Vestergaard (2025). "Large language models, small labor market effects." NBER Working Paper 33777, National Bureau of Economic Research.

Lee, David S. (2009). "Training, wages, and sample selection: Estimating sharp bounds on treatment effects." *The Review of Economic Studies*, 76(3), pp. 1071–1102.

Tucker, Lee (2026). "You're (not) hired: Artificial intelligence and early career hiring in the quarterly workforce indicators." CES Working Paper 26-27, U.S. Census Bureau, Center for Economic Studies.

Webb, Michael (2020). "The impact of artificial intelligence on the labor market." Working paper. SSRN 3482150.

Appendix A AI Exposure Measures

Table A1: Occupation-level AI exposure measures

Source (shorthand)	What it is intended to capture	How it is constructed
Felten et al. (2021) *(AIOE)*	Potential exposure of an occupation to progress in AI broadly, working through the human abilities the occupation requires. A capability-based measure of exposure, agnostic about whether the effect is substituting or complementary.	Ten AI applications (from the Electronic Frontier Foundation's AI Progress Measurement benchmarks) are linked to 52 O*NET abilities via a crowdsourced application–ability relatedness matrix. An ability-level exposure is formed as the (equally weighted) sum of the ten relatedness scores, then aggregated to the occupation by weighting each ability by its O*NET prevalence and importance. Continuous AIOE score by 6-digit SOC (774 occupations).
Felten et al. (2023) *(LM-AIOE)*	The same ability-based exposure concept, but reweighted to capture exposure specifically to *language modeling* (ChatGPT-type generative language AI) rather than to AI in general.	Identical AIOE framework and the same application–ability relatedness matrix, but the application weights are changed from equal (all ten set to 1) to 1 for the single "language modeling" application and 0 for the other nine. The ability-level exposure therefore counts only abilities related to language modeling; occupation aggregation (prevalence- and importance-weighted) is unchanged. 6-digit SOC (774 occupations).
Eloundou et al. (2024) *(Eloundou β)*	The share of an occupation's work tasks for which access to LLMs (and LLM-powered software built on them) could substantially cut the time to complete the task—a task-based measure of generative-AI capability exposure.	Each O*NET task is rated based on the ability of direct LLM access *or* additional LLM-powered software to reduce task time by $\geq$50%. Occupation exposure is the exposed task share. Paper provides ratings done by GPT-4 and by human; focus here on the GPT-4 ratings and the β formulation of the exposed task share, which gives less weight to tasks that require additional LLM-powered software to realize task time gains. 923 O*NET-SOC occupations.

Continued on next page

PRELIMINARY AND INCOMPLETE: DO NOT CITE OR CIRCULATE

Table A1, continued

Source (shorthand)	What it is intended to capture	How it is constructed
Frey and Osborne (2017) *(Frey–Osborne)*	The probability that an occupation could be fully automated given foreseeable advances in machine learning and mobile robotics. A pre-LLM automation-susceptibility benchmark.	Seventy occupations are hand-labeled automatable/not by ML experts; nine O*NET variables capturing three "engineering bottlenecks" (perception and manipulation, creative intelligence, social intelligence) are used to predict computerization probabilities for other occupations via a Gaussian-process classifier (702 occupations, 6-digit 2010 SOC).
Handa et al. (2025) *(Anthropic AEI)*	Realized, usage-based exposure: how much AI is actually being used to perform an occupation's work today, rather than how much it could be in principle.	A large sample of real Claude.ai conversations is mapped to the O*NET tasks they correspond to, and then to occupations. Occupation exposure is the share of the occupation's O*NET tasks observed being performed in Claude conversations (756 occupations, current SOC).
Eisfeldt et al. (2025) *(ESTZ GenAI)*	The share of an occupation's tasks that current generative AI can make more productive—a generative-AI capability exposure built independently of Eloundou β.	GPT-4 classifies each of ∼19,000 O*NET tasks (as of March 2023) for whether GPT-4-level tools raise productivity on the task; occupation exposure is the resulting task share. Total exposure decomposes into core + supplemental tasks; here consider *total* and *core* measures (840 occupations, 6-digit 2010 SOC).
Brynjolfsson et al. (2018) *(BMR SML)*	Suitability for machine learning—how amenable an occupation's work is to being performed by supervised ML. A pre-generative-AI task-automation benchmark, distinct from Frey–Osborne's full-job automation probability (it scores activities, not whole jobs).	Crowd raters score each O*NET Detailed Work Activity on a 21-item ML-suitability rubric; the DWA scores are spread to all O*NET tasks and importance-weighted to occupations. Higher = more ML-suitable (771 occupations, 6-digit 2010 SOC, from O*NET release 21.1).

Continued on next page

PRELIMINARY AND INCOMPLETE: DO NOT CITE OR CIRCULATE

Table A1, continued

Source (shorthand)	What it is intended to capture	How it is constructed
Webb (2020) *(Webb AI)*	Exposure to AI as revealed by what AI has actually been invented to do: an occupation is exposed to the extent that the tasks it is made of are the tasks AI patents describe performing.	Verb–object pairs are extracted from the text of AI patents and, separately, from O*NET task descriptions. An occupation's score is the extent to which the verb–object pairs in its tasks appear among those in AI patents, aggregated across the occupation's tasks and expressed as a percentile across occupations. Webb constructs parallel scores for software and industrial robots; the focus here is on the AI series (`pct_ai`). Published on `occ1990dd` codes, here mapped to 2018 SOC (768 scored occupations).

PRELIMINARY AND INCOMPLETE: DO NOT CITE OR CIRCULATE

Table A2: Comparing AI Exposure Measures

(a) Cross-occupation correlations among exposure measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) AIOE	1							
(2) LM-AIOE	0.98	1						
(3) Eloundou β	0.86	0.84	1					
(4) ESTZ GenAI	0.80	0.79	0.89	1				
(5) Anthropic AEI	0.57	0.58	0.65	0.68	1			
(6) BMR SML	0.30	0.28	0.33	0.30	0.26	1		
(7) Frey–Osborne	−0.44	−0.46	−0.28	−0.20	−0.15	0.08	1	
(8) Webb AI	0.07	0.03	0.08	0.08	0.02	0.02	−0.13	1

(b) Cross-occupation rank correlations among exposure measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) AIOE	1							
(2) LM-AIOE	0.97	1						
(3) Eloundou β	0.87	0.86	1					
(4) ESTZ GenAI	0.84	0.83	0.91	1				
(5) Anthropic AEI	0.70	0.71	0.74	0.73	1			
(6) BMR SML	0.29	0.29	0.34	0.32	0.32	1		
(7) Frey–Osborne	−0.37	−0.39	−0.23	−0.22	−0.23	0.09	1	
(8) Webb AI	−0.01	−0.06	0.02	0.01	0.02	−0.04	−0.10	1

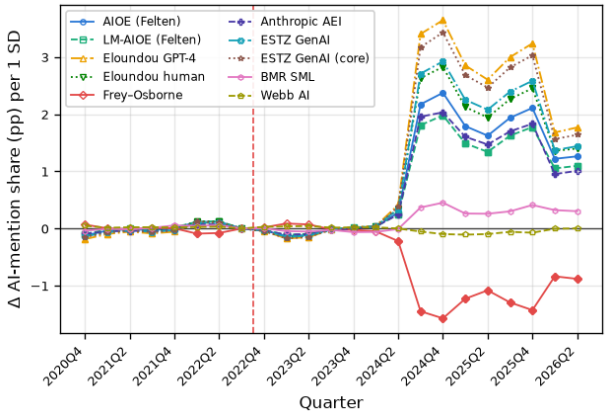
(c) Ten highest-exposure occupations under a generative-AI measure, Frey–Osborne, BMR, and Webb

	LM-AIOE (generative AI)	**Frey–Osborne p(automation)**	**BMR SML (ML suitability)**	**Webb AI (patent text)**
1	Telemarketers	Telemarketers	Mechanical Drafters	Buyers and Purchasing Agents, Farm Products
2	English Language and Literature Teachers, Postsecondary	Mathematical Technicians	Concierges	Chemical Engineers
3	Foreign Language and Literature Teachers, Postsecondary	Title Examiners, Abstractors, and Searchers	Morticians, Undertakers, and Funeral Directors	Materials Engineers
4	History Teachers, Postsecondary	Cargo and Freight Agents	Brokerage Clerks	Atmospheric and Space Scientists
5	Law Teachers, Postsecondary	Insurance Underwriters	Psychologists, All Other	Optometrists
6	Philosophy and Religion Teachers, Postsecondary	Watch Repairers	Gaming Cage Workers	Pest Control Workers
7	Sociology Teachers, Postsecondary	Library Technicians	Multi-Media Artists and Animators	Construction and Building Inspectors
8	Political Science Teachers, Postsecondary	New Accounts Clerks	Ushers, Lobby Attendants, and Ticket Takers	Elevator and Escalator Installers and Repairers
9	Criminal Justice and Law Enforcement Teachers, Postsecondary	Photographic Process Workers and Processing Machine Operators	Office Machine Operators, Except Computer	Water and Wastewater Treatment Plant and System Operators
10	Sociologists	Tax Preparers	File Clerks	Inspectors, Testers, Sorters, Samplers, and Weighers

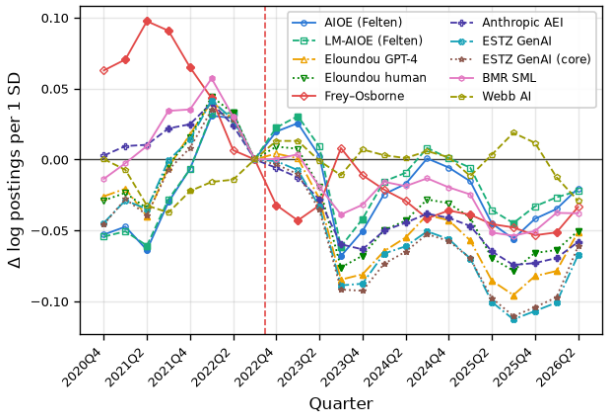
Notes: Panel (a) reports Pearson correlations and panel (b) Spearman rank correlations across the 456 five-digit SOC occupations the measures span, computed pairwise (each cell uses the occupations the two measures share). Panel (c) lists the ten occupations that select measures score highest, at that measure's native resolution and using the source authors' occupation titles.

Figure A1: Lightcast Event Study Estimates, All Exposure Measures

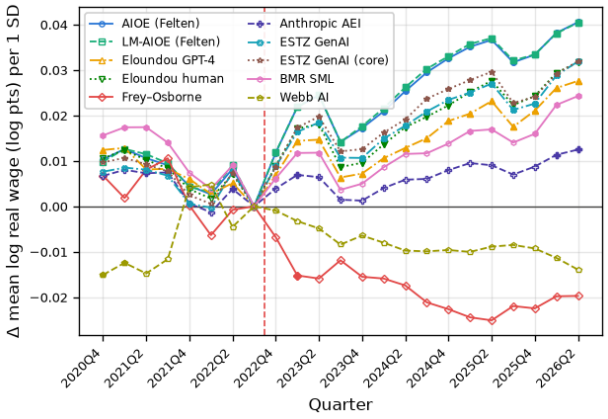
(a) Percent mentioning AI-related terms



(b) log(job ads)



(c) log(real posted wage)



Notes: Coefficients β_t from equation 1, estimated separately for each exposure measure, with 2022Q3 omitted and exposure measured in standard deviations. Confidence intervals are omitted for legibility; markers are filled where the estimate differs from zero at the 5 percent level. Balanced panels contain 619 to 641 occupations depending on the measure. Standard errors clustered on occupation.

Appendix B Text Analysis

This section describes the methodology used in the text analysis presented in Figure B3. My approach produces three measures aimed at capturing differences between AI ads and non-AI ads, excluding direct mentions of AI, in terms of simple text (excess cosine dissimilarity), the ability of the text to predict AI status, and the work activities they reflect.

Appendix B.1 Tagging AI-related Terms in Job Ads

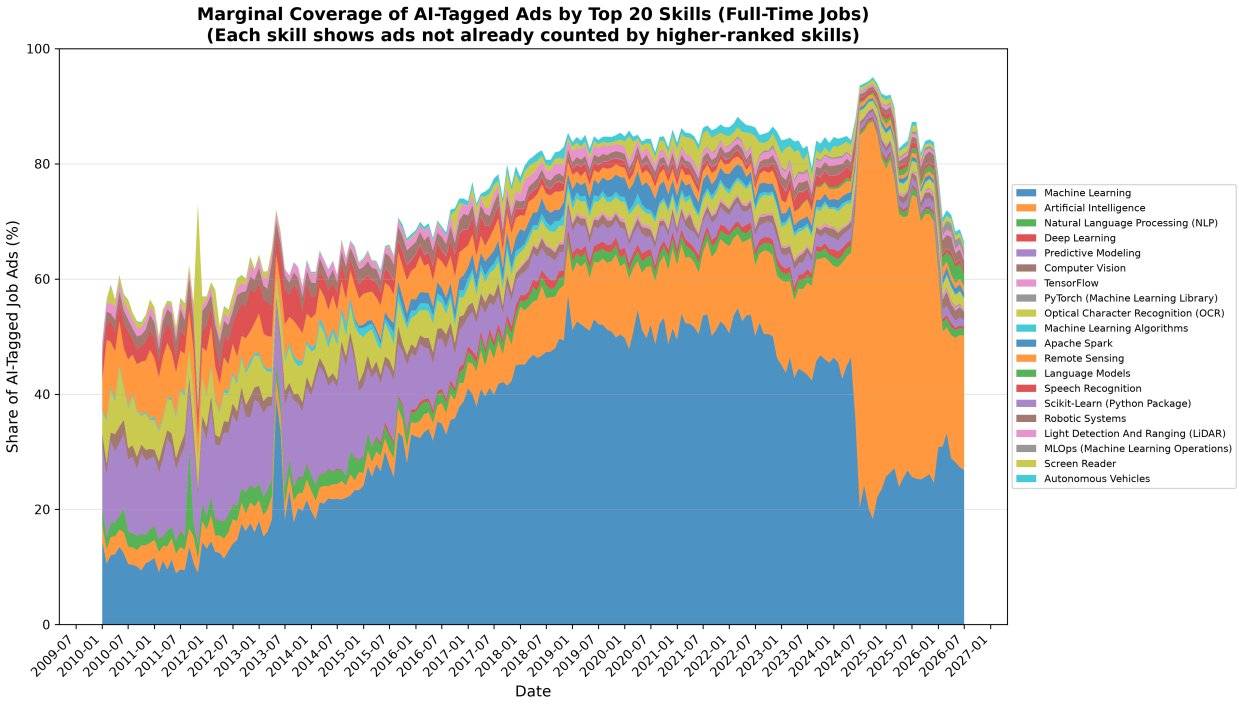
I search the text of job ads obtained from Lightcast for mentions terms related to AI. The core of this search is based on a set of 350 AI-related skills, provided by Lightcast. For terms that are often used interchangeably with a well-known acronym (e.g. NLP, OCR, etc.) I also search for the acronym.

The most commonly used terms capture a large majority of the ads that mention any AI-related terms. In 2026Q2, the top two terms, machine learning and articial intelligence, capture more than 40 percent of ads that mention any AI-related terms, as shown in Figure B1. The top 20 terms capture nearly 70 percent of ads in that quarter and over 90 percent of ads at their peak.

Not every mention of an AI-related term provides information about the content of the job, and simply counting mentions of these terms can be misleading. In particular, I attempt to correct for three issues: terms that have pre-existing meanings that are unrelated to AI, terms used in employer branding or other contexts that do not pertain to the content of the job being advertised, and terms used in the context of disclosing AI use in hiring processes.

To address potential ambiguity around terms that have both AI-related and non-AI meanings (e.g. “transformer” could also refer to an electrical transformer, “torch” could also refer to the welding tool, etc.), I do not count as mentioning AI any ad that mentions only an ambiguous term, unless the ambiguous term is mentioned in the same sentence as at least one term related to machine learning (e.g. language model, deep learning, data science,

Figure B1: Marginal share of AI-mentioning job ads captured by top terms



Notes: Each band is the share of AI-mentioning full-time ads that a term captures beyond all higher-ranked terms, so the bands do not double count and their total is the share of AI-mentioning ads reached by the twenty most common terms. Terms are ranked by their share over the full period.

various Python package names, etc.).

To identify mentions of AI that are part of employer branding or otherwise used in ways that are unlikely to reflect the content of the job being advertised, I search for "AI" used in combination with various fairly generic words (e.g. AI-powered, advanced AI, predictive AI, AI company, AI products, leverage AI, etc.).

To identify mentions of AI-related terms that occur in the context of disclosing AI use in hiring processes, I look within the sentence containing each mention of an AI-related term (up to 250 characters before and after the mention) for terms related to hiring processes: hiring, recruit, applicant, candidate, interview, talent acquisition, employment decision, selection process, job application, fair chance, background check, lie detector, polygraph, e-verify. Ads that mention one or more AI-related terms but only within sentences that also include hiring-process terms are classified as making disclosure-only mentions of AI; these are shaded in purple in Figure 1a.

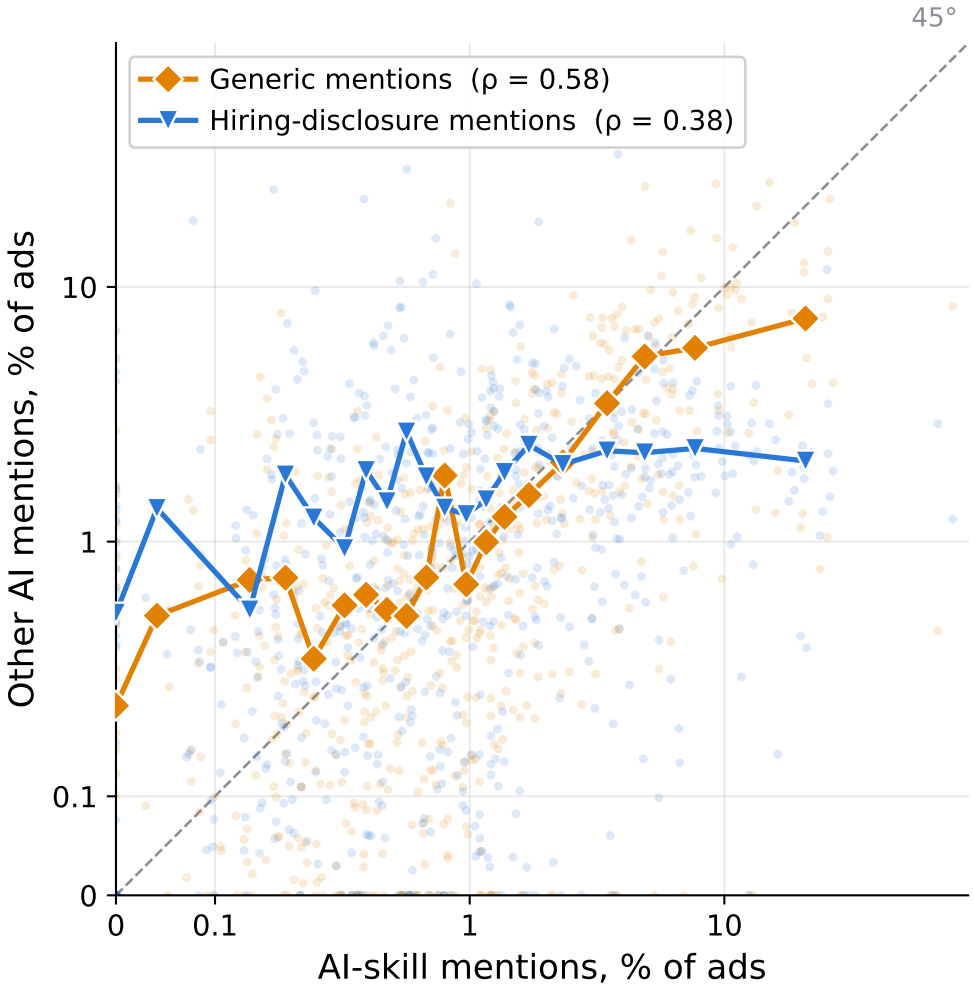
At the end of the period considered, both hiring process disclosure mentions of AI and more generic/branding-related mentions are positively correlated with mentions of AI skills at the occupation level (see Figure B2). This suggests that naive measurement could overstate the degree to which employers are mentioning AI skills in job ads.

After tagging ads' AI-mention status, I save a version of the text of each AI-mentioning ad that has had AI-related terms stripped from it. This version is used in subsequent comparisons across types of ads, in order to focus on differences in text across groups beyond that generated by defining the groups based on differences in text (i.e. mentions of AI-related terms).

Appendix B.2 Cosine Similarity Analysis

To compare AI and non-AI ads using cosime similarity, I convert the AI-masked text of ads into vectors of terms, with each term consisting of one or two words. I exclude terms that appear exactly once across ads and those that appear in more than 90 percent of ads. I

Figure B2: Occupational AI mention rates by type of mention, 2026Q2



Notes: Each faint point is a five-digit occupation in 2026Q2 with at least 125 full-time ads. The horizontal axis is the share of the occupation's ads naming a curated AI skill; the vertical axis is the share whose only AI content is generic "AI" language or hiring-disclosure language. Heavy markers are means within twenty equal-count bins of the horizontal axis, which the two series share. Axes are on a symmetric log scale so that occupations at zero remain visible, and the diagonal is the 45-degree line.

apply term frequency-inverse document frequency (TF-IDF) weights to the remaining terms such that for term t:

$$\begin{aligned} \text{Weighted Frequency} &= \text{Frequency} \times \text{TF Weight} \times \text{IDF Weight} \\ & n \times (1 + ln\,(n)) \times \left(ln \left(\frac{1+n}{1+df\,(t)} \right) + 1 \right) \end{aligned} \tag{B1}$$

where n is the unweighted number of times term t appears in an ad and $df(t)$ is the number of times term t appears across all ads in the comparison set (in this case, an occupation-quarter cell). The weighted vector is then normalized to have length = 1.

I compare ads with and without mentions of AI-related terms within five-digit SOC occupation-quarter cells. I focus on cells with at least 50 ads of each type. I sample up to 500 AI ads and up to 2000 non-AI ads within each occupation-quarter and take the cosine of the vectors for the ads within each cross-type pair of ads within that sample. For comparison, I also take the cosine between the vectors for all distinct pairs of non-AI ads. I then average these pairwise cosine values within cells to produce average similarity of AI ads to non-AI ads and average similarity of non-AI ads to each other. I difference these to produce a measure of excess dissimilarity within each occupation-quarter cell:

$$\text{Excess Dissimilarity} = \overline{cos\,(n, n')} - \overline{cos\,(a, n)} \tag{B2}$$

where n and n' represent all pairs of distinct non-AI ads and a and n represent all pairs of AI and non-AI ads. The overall excess dissimilarity measure, which is plotted in Figure B3, is an aggregation of cell-level measures weighted by the number of adds that mention AI-related terms.

Appendix B.3 Classifier Discriminability Analysis

I also produce a measure of how distinctive the non-AI portions of AI-mentioning ads are within their occupations. Operating within occupation-quarter cells that contain at least 25

ads that mention AI-related terms, I train a classifier to predict AI status based on the same AI-masked, TF-IDF weighted text of job ads used in the cosine similarity analysis. Each cell is split into five equal segments, and predictions for each segment are based on training conducted using the other four.

Using each ad's predicted probability of being an AI-mentioning ad, I then compare every possible pair of one AI ad with one non-AI ad within the cell to determine the share for which the classifier estimated a higher probability of being an AI ad for the ad that actually is an AI ad. This share is equivalent to the area under the curve (AUC) for the classifier's receiver operating characteristic (ROC) curve and serves as a measure of the discernabiltiy of AI ads from non-AI ads. As with the cosine similarity analysis above, this measure excludes the influence of actual AI-related terms by construction.

The series plotted in Figure B3 is again an aggregation of cell-level AUC measures. As in the cosine similarity analysis, cells are weighted by the number of AI ads they contain.

Appendix B.4 O*NET Task Content Analysis

Finally, I consider a measure of how the job-related content of AI and non-AI ads aligns with a known set of work activites obtained from O*NET. I use a small sentence-embedding model (all-MiniLM-L6-v2) accessed through Hugging Face to vectorize O*NET work activity statements (for both general and detailed work activities) and each sentence in the text of each job ad. To limit computational requirements, I use a 10 percent sample of job ads and begin the analysis in 2020Q4, two years prior to the introduction of ChatGPT.

Each sentence from each ad is then matched to its closest work activity, based on cosine similarity of their vectors. Job ad sentences that are not sufficiently close to any work activity are skipped. Once activities have been assigned to sentences, I construct a vector giving the distribution of work activities in each ad. This vector has one entry for each work activity (there are 41 general work activities and 2,087 detailed work activities). Each entry in the vector is the number of sentences matched to a work activity divided by the total number

of sentences matched to any work activity.

Within occupation-quarter cells with at least 5 AI ads and 5 non-AI ads, I then average the work activity distribution vectors within the sets of AI and non-AI ads. I then calculate the Jensen-Shannon (JS) divergence betwenn these two average vectors:

$$JSD\left(A||N\right) = \frac{1}{2}D\left(A||M\right) + \frac{1}{2}D\left(N||M\right) \tag{B3}$$

where A is the average distribution of work activities in AI ads, N is the average distribution of work activities in non-AI ads, M is the average of A and N, and D is the Kullback-Leibler divergence between each ad type's vector and their average. In general, the Kullback-Leibler divergence between two probability distributions P and Q is given by

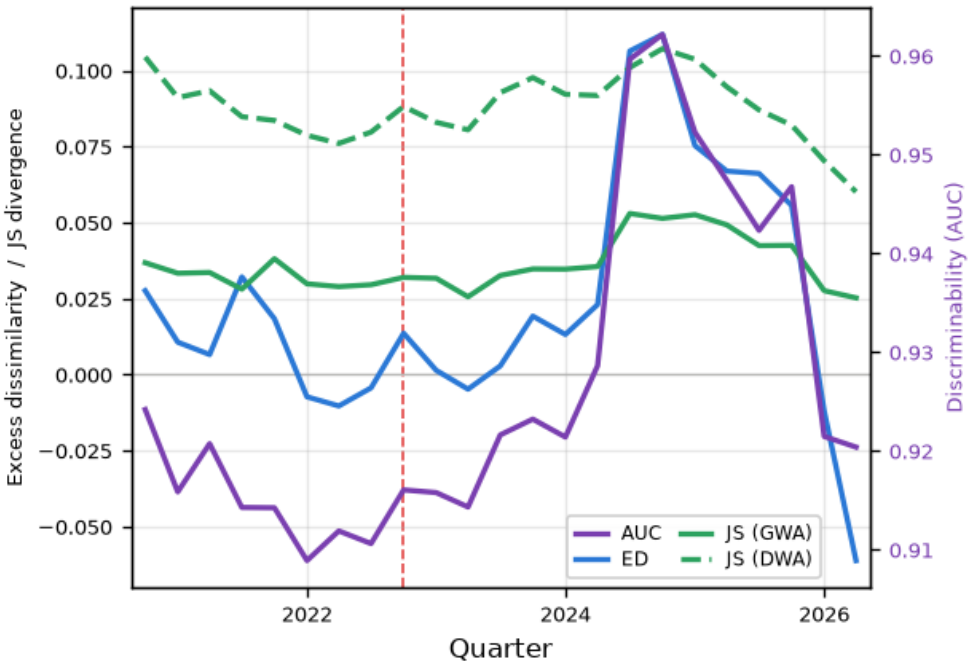
$$D_{KL}(P,Q) = \sum_{x \in X} P(x) log \frac{P(x)}{Q(x)}. \tag{B4}$$

I also split non-AI ads into two halves at random and calculate the JS divergence between the two halves as a measure of baseline differences between ads in each cell.

As in the prior analyses, the national estimates reported in Figure B3 are aggregations of cell-level measures constructed by averaging across cells with the number of AI ads in each cell as weights. Figure B3 reports excess JS divergence between AI ads and non-AI ads by subtracting the average placebo distance between random sets of non-AI ads from the average distance between AI ads and non-AI ads.

Each of these measures shows the content of AI and non-AI ads diverging temporarily around the spike in AI mentions in job ads that reverts by the end of the period considered.

Figure B3: Text similarity



Notes: National quarterly aggregates of the three cell-level measures described in this appendix, comparing AI-mentioning ads with other ads in the same five-digit occupation after masking the AI terms themselves. Excess dissimilarity (ED) and the Jensen–Shannon divergences over general (GWA) and detailed (DWA) work activities are read on the left axis, each net of a placebo computed between random halves of the non-AI ads; classifier discriminability (AUC) is read on the right axis, where 0.5 would mean the two groups are indistinguishable. Cells are weighted by their number of AI-mentioning ads. The dashed line marks the release of ChatGPT.

Appendix C Selection into Wage Posting

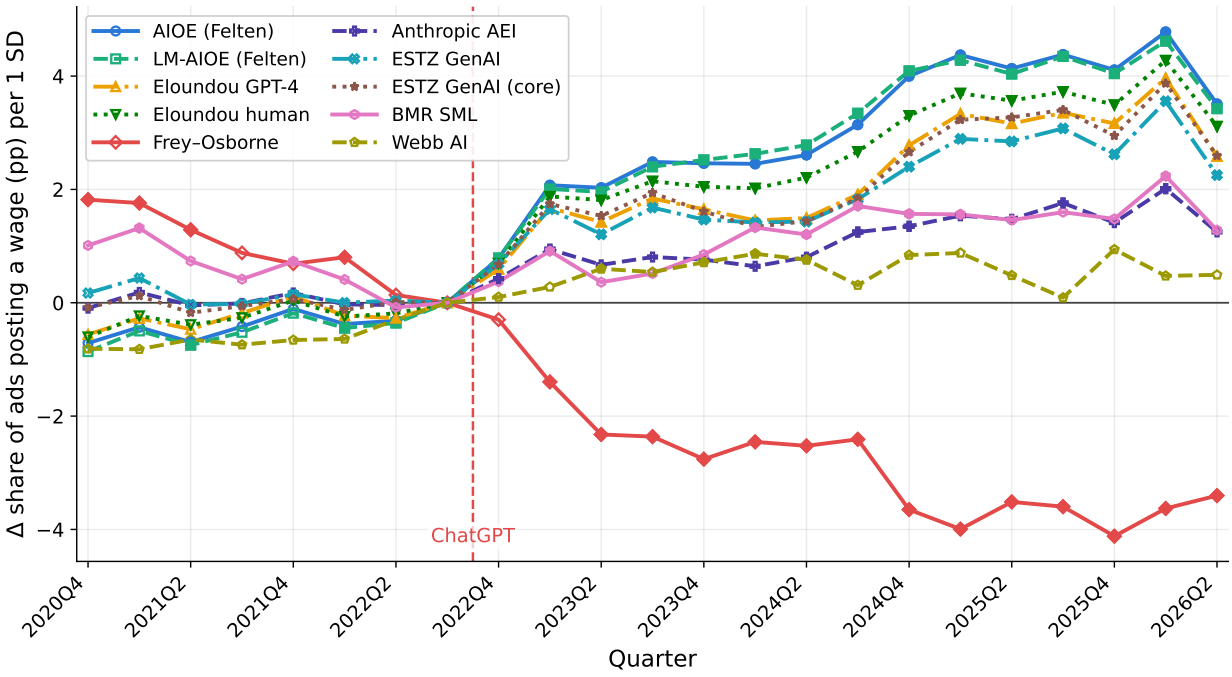
It has become much more common for ads to include wage information, especially over the last five years or so. That trajectory is also evident in all ads, regardless of AI-related content. Shifts of this magnitude raise questions about the composition of firms that post wages and the comparability of posted wage measures over time, given potential changes in composition.

Compositional concerns are heightened by the fact that wage posting has become more common particularly in more AI-exposed occupations since the release of ChatGPT. Figure C1 shows systematically higher wage posting rates in occupations with higher levels of LLM/GenAI-related exposure, beginning at the end of 2022. Increases in wage posting over the period of interest are not only substantial, but also clearly correlated with exposure to AI. This could make comparisons of aggregate wage information over time deceiving if, for example, firms or jobs that unobservably make more productive use of AI and pay higher wages are also more likely to include wage information in job ads now than they used to be.

Unrelated changes in policy requiring wage information to be included in job ads in some places and for some types of firms further complicates analysis of posted wages. Table C1 describes state-level laws passed recently that require firms to post wage information in job ads, while Table C2 describes other laws related to pay transparency that either were passed by sub-state jurisdictions or require wage information to be provided in some way other than in job ads.

Conducting an event study analysis of the state-level policies listed in Table C1 that were in effect during the period considered shows that the rate at which employers include wage information in job ads increases substantially when they are requried to do so. Figure C2 shows that, following a modest trend up in wage posting ahead of these laws taking effect (which could reflect anticipatory compliance following enactment), wage posting rates jump more than 10 percentage points in the quarter the law takes effect. Wage posting rates continue to rise in subsequent quarters, reaching roughly 20 percentage points above the

Figure C1: Relationship between AI exposure and wage posting, by exposure measure



Notes: Coefficients β_t from equation 1 with the share of an occupation's ads that post a wage as the outcome, in percentage points, estimated separately for each exposure measure. 2022Q3 is omitted and exposure is measured in standard deviations; markers are filled where the estimate differs from zero at the 5 percent level. Standard errors clustered on occupation.

Table C1: Pay-transparency posting requirements

State	Effective	Quarter	Coverage
Panel (a): In force within the estimation window			
Colorado	Jan. 1, 2021	2021Q1	All employers
California	Jan. 1, 2023	2023Q1	15+ employees
Washington	Jan. 1, 2023	2023Q1	15+ employees‡
New York	Sep. 17, 2023	2023Q3	4+ employees‡
Hawaii	Jan. 1, 2024	2024Q1	50+ employees‡
District of Columbia	Jun. 30, 2024	2024Q2	All employers
Maryland	Oct. 1, 2024	2024Q4	All employers
Illinois	Jan. 1, 2025	2025Q1	15+ employees
Minnesota	Jan. 1, 2025	2025Q1	30+ employees†
New Jersey	Jun. 1, 2025	2025Q2	10+ employees
Vermont	Jul. 1, 2025	2025Q3	5+ employees
Massachusetts	Oct. 29, 2025	2025Q4	25+ employees†
Panel (b): Effective after the estimation window			
Virginia	Jul. 1, 2026	2026Q3	All employers
Maine	Jul. 29, 2026	2026Q3	10+ employees‡
Connecticut	Oct. 1, 2026	2026Q4	All employers
Delaware	Sep. 26, 2027	2027Q3	More than 25 employees‡

Notes: State laws requiring employers to include a pay range in the job posting itself. An ad is treated as covered from the quarter containing the effective date onward, keyed on the state in which the job is posted. Panel (a) lists the twelve states in force within the estimation window (through 2026Q2) and used in the analysis; panel (b) lists requirements taking effect after it. Coverage is the employer-size threshold at which the requirement binds, where "all employers" means any employer with at least one employee in the jurisdiction. Thresholds marked † count in-state employees only and those marked ‡ are left unspecified by the statute, so the column is not strictly comparable across rows; the treatment indicator itself applies no size screen. Substate ordinances, and laws requiring disclosure only on request, are listed in Table C2. Effective dates are from industry pay-transparency law trackers and coverage from the statute or the enforcing state agency, accessed July and August 2026.

Table C2: Pay-transparency laws excluded from the mandate indicator

Jurisdiction	Effective	Requirement and reason for exclusion
Panel (a): Disclosure on request or before offer, not in the posting		
Maryland	2020	Range on applicant request. Superseded by Maryland's Oct. 1, 2024 posting requirement, which is used
Connecticut	2021	Range on applicant request. Superseded by Connecticut's Oct. 1, 2026 posting requirement, which falls after the window
Nevada	Oct. 2021	Range provided after an interview
Rhode Island	Jan. 2023	Range on applicant request
Panel (b): Substate ordinances the state-level indicator cannot represent		
Jersey City, NJ	Apr. 2022	Posting requirement; precedes New Jersey's Jun. 1, 2025 law
Ithaca, NY	Sep. 1, 2022	Posting requirement; precedes New York's Sep. 17, 2023 law
New York City, NY	Nov. 1, 2022	Posting requirement; precedes New York's Sep. 17, 2023 law
Westchester County, NY	Nov. 6, 2022	Posting requirement; precedes New York's Sep. 17, 2023 law
Cleveland, OH	Oct. 2025	Posting requirement; Ohio has no state posting law

Notes: Pay-transparency laws in force during the estimation window that are not treated as posting mandates. The laws in panel (a) require an employer to supply a pay range on request, at interview, or before an offer, but not to publish one in the advertisement, so they do not mechanically raise the share of ads that disclose a wage. The ordinances in panel (b) are posting requirements, but the treatment indicator is keyed on the state and cannot represent substate variation; each predates its own state's law, or has no state analogue, so state-level coverage understates wage posting in these labor markets before the state effective date. This matters most for New York, where New York City was covered roughly ten months before the state. Sources as in Table C1.

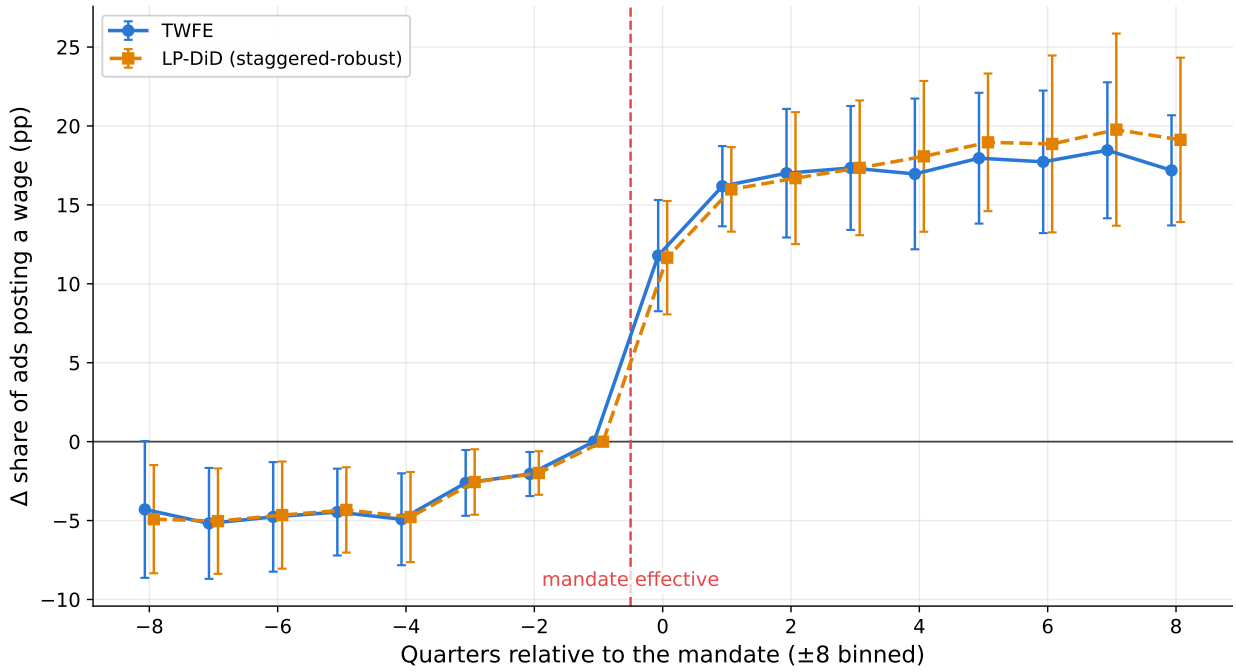
reference period and about 25 percentage points above the trend that prevailed a year before the laws took effect. Effects are similar whether estimated using a traditional two-way fixed effects approach or the Dube et al. (2025) local projections approach.

To address potential concerns about how changing selection into wage posting could affect wage estimates, Figure C3 in the main text presents estimates from alternative specifications intended to limit the influence of selection in various ways:

- Incumbent pairs: produce posted wage estimates using a sample that includes only employer-occupation pairs that are observed with posted wages both before and after the introduction of ChatGPT, reducing the scope for selection into wage posting on the extensive margin to affect estimates

- Employer-occupation fixed effects: do not restrict the sample, but include employer-occupation fixed effects in the estimation, so that extensive margin selection into wage posting does not contribute to identification of the relationship between AI exposure

Figure C2: Wage posting rates around the introduction of state-level pay transparency laws



Notes: Effect of state pay-transparency posting mandates on the share of ads that post a wage, estimated on a state-quarter panel with state and quarter fixed effects and event time measured from the quarter a mandate takes effect. Estimates are shown for a two-way fixed effects specification and for the local projections estimator of Dube et al. (2025), which is robust to mandates taking effect in different quarters across states. Treated states are listed in Table C1. Standard errors clustered on state; 95 percent confidence intervals.

and posted wages

- DiNardo et al. (1996) reweighting: Reweight ads within occupation-quarter cells to match the pre-ChatGPT period's joint distribution of state and education requirements among ads that posted wages, countering the influence of any selection into wage posting on these observable dimensions

- Hainmueller (2012) entropy balancing: Reweight occupation-quarter cells to match the pre-ChatGPT means of education and experience requirement and wage posting law measures among ads that posted wages, keeping weights as close to uniform as possible, another approach to countering the influence of selection on observables

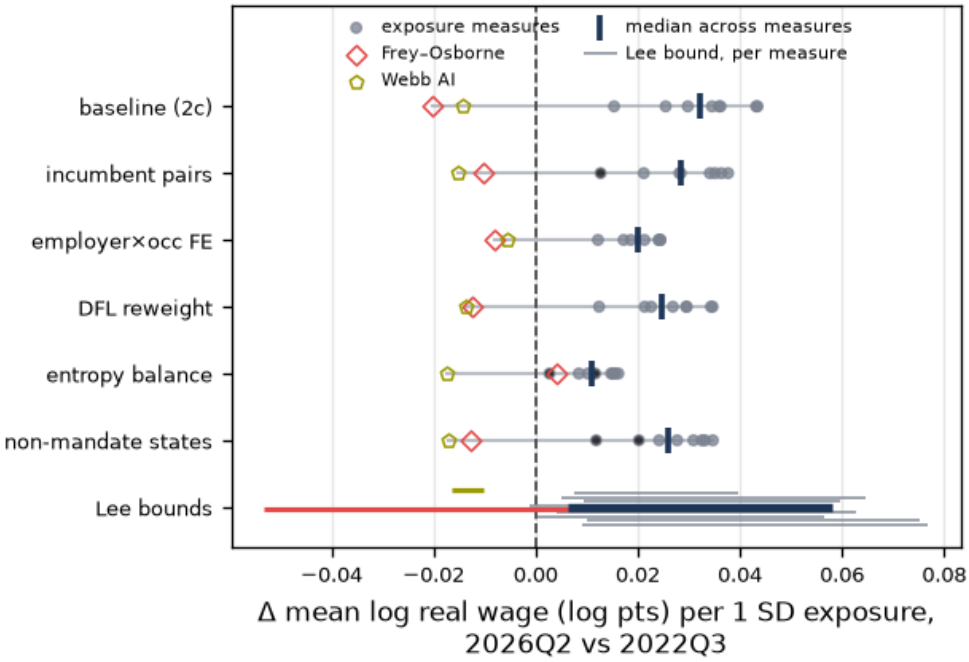
- Non-mandate states: produce posted wage estimates using only states that did not have a wage posting law in effect during the period considered, sidestepping the large increase in wage posting associated with such laws

- Lee (2009) bounds: Trim each occupation-quarter cell's "excess" wage posting from the top/bottom of the distribution to match trend growth in wage posting for the median occupation, then re-estimate effects on posted wages.

Figure C3 summarizes how these approaches alter the wage estimate for 2026Q2. In all cases, more LLM-exposed occupations see higher real posted wages in that quarter. In most cases, the Lee bounds lie entirely above zero.

While a Heckman (1979) correction may seem like a natural option for addressing selection in this context, especially since pay disclosure laws shift rates of wage posting, it seems less suitable upon closer examination. Wage posting laws seem unlikely to satisfy the exclusion restriction. Figure C4 shows estimated effects of the introduction of state-level wage posting laws on posted wages, estimated using a stacked difference-in-differences approach with clean controls analogous to the approach taken in Cengiz et al. (2019). The figure shows effects estimated using all ads, which appear to be flat though somewhat imprecise,

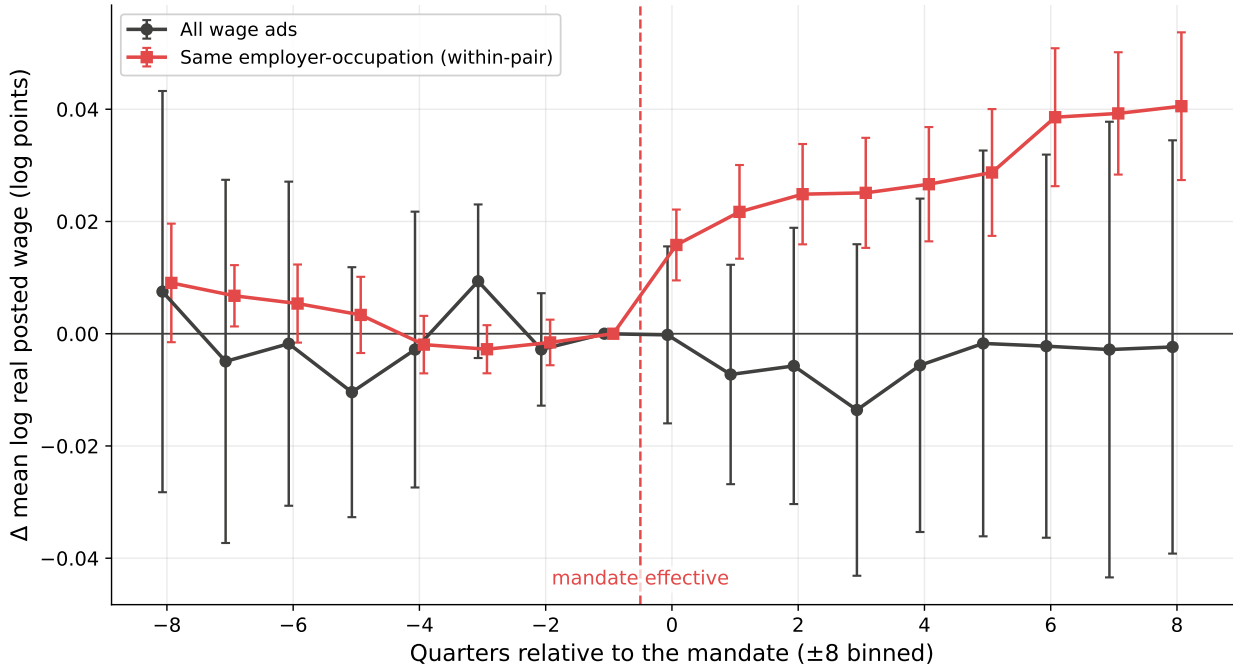
Figure C3: Wage selection adjustments



Notes: Change in the occupation-level mean of log real posted wages between 2022Q3 and 2026Q2 per standard deviation of AI exposure, re-estimated under each correction for selection into wage posting described above. Each gray point is one exposure measure and the vertical bar is the median across measures; Frey–Osborne and Webb, the two measures not based on large language models, are marked separately. The bottom row reports Lee (2009) bounds: thin lines are the bounded interval for each measure and the thick bar spans the median lower and upper bounds.

and effects estimated using only employer-occupation pairs that were posting wages before the wage posting laws took effect, which shows a clear, immediate increase in posted wages after these laws take effect. The magnitude increases to about 4 percent by two years later. Together, these estimates highlight the importance of the selection concerns and illustrate that wage posting laws affect the wage level directly and therefore do not satisfy the exclusion restriction required in a Heckman-style analysis.

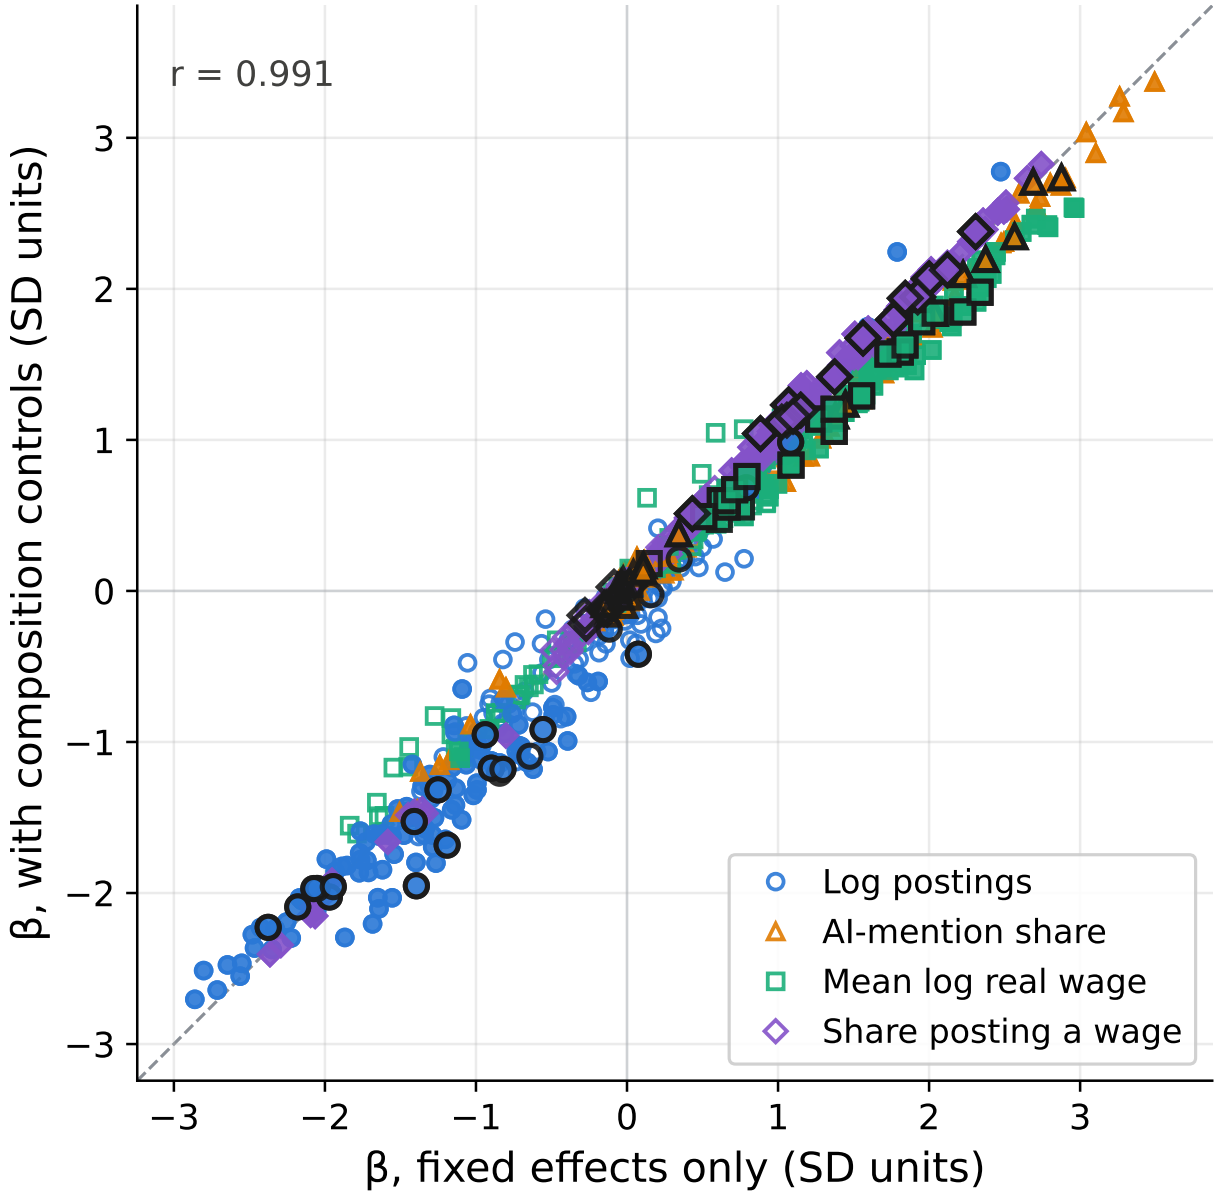
Figure C4: Posted wages around the introduction of state-level pay transparency laws



Notes: Effect of state pay-transparency posting mandates on the level of posted wages, from a stacked difference-in-differences design in which each mandate is compared with states not yet mandated over a clean event window. The first series uses all ads posting a wage; the second restricts to employer-occupation pairs already posting wages before the mandate, so that it holds the set of wage posters fixed and the difference between the two series reflects the changing composition of ads that disclose pay. Standard errors clustered on state; 95 percent confidence intervals.

Appendix D Additional Tables and Figures

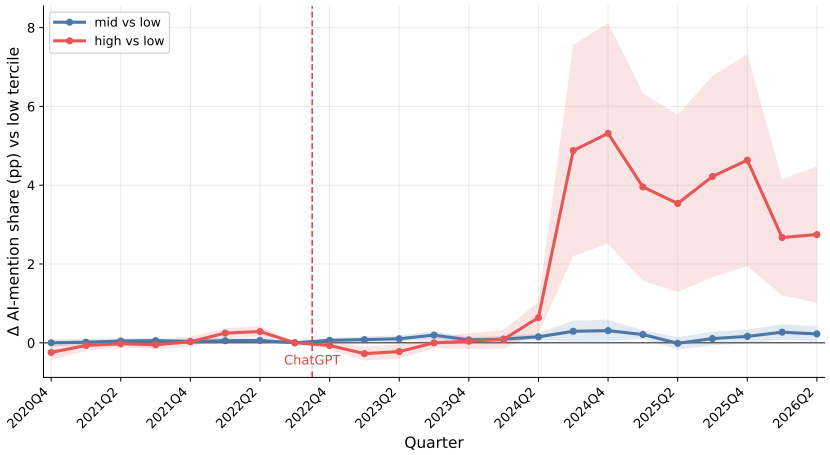
Figure D1: Event study coefficients with and without controls



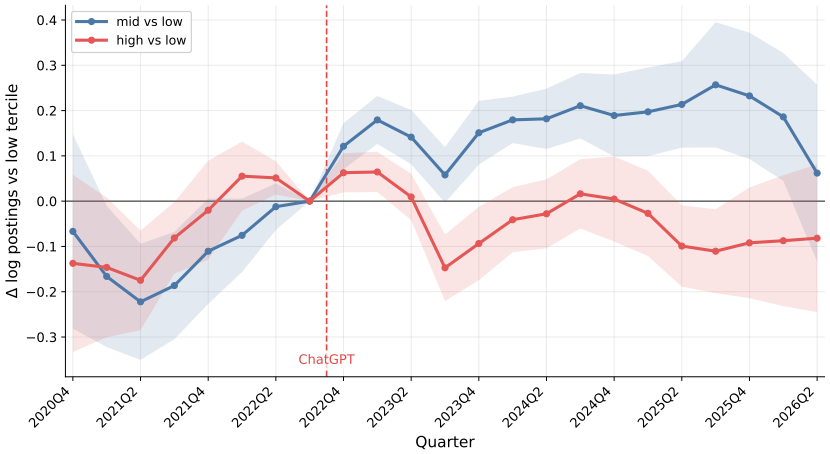
Notes: Each point is one event-study coefficient—an exposure measure in a given quarter—estimated with controls for the education and experience requirements stated in the ad, plotted against the same coefficient estimated without them. Coefficients are divided by the standard deviation of the uncontrolled estimates for their outcome so that outcomes measured in different units share an axis. Points on the 45-degree line indicate that the controls leave the estimate unchanged. Ringed markers are the principal-component measure.

Figure D2: Event study estimates, dose-response specification

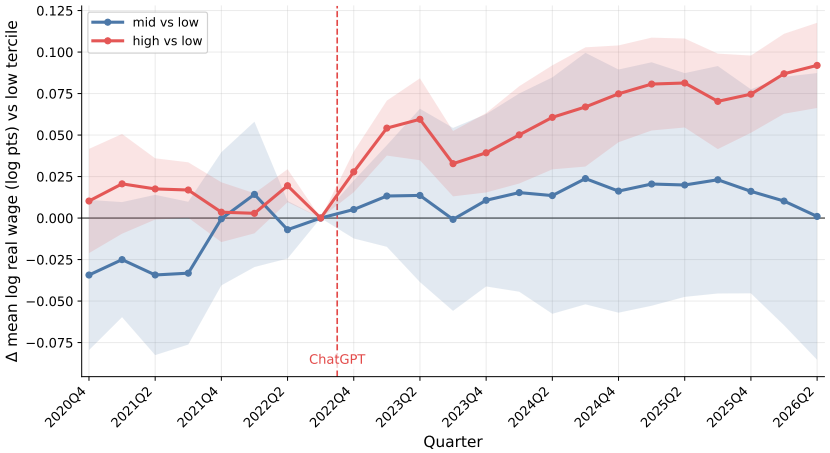
(a) AI share



(b) log(job ads)

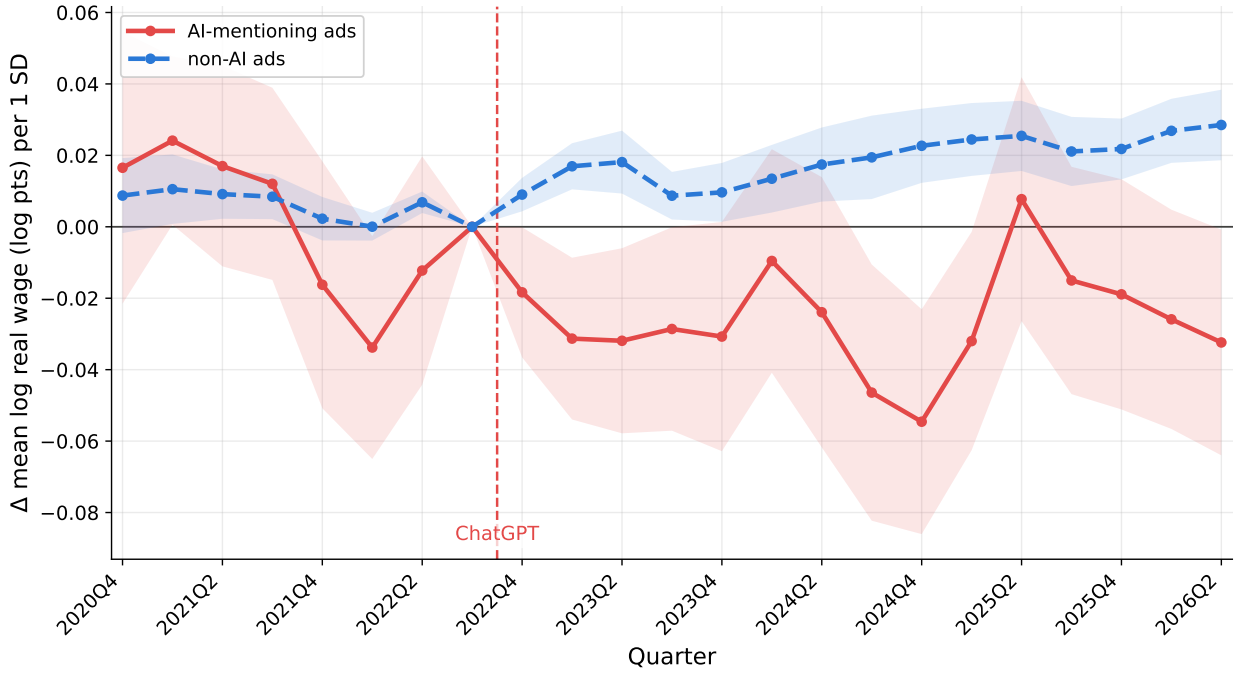


(c) Posted wages



Notes: Occupations are sorted into terciles of AI exposure, measured by the first principal component of the exposure measures based on large language models, and each panel plots the difference between the middle and highest terciles and the lowest tercile by quarter, relative to 2022Q3. Shaded bands are 95 percent confidence intervals from standard errors clustered on occupation.

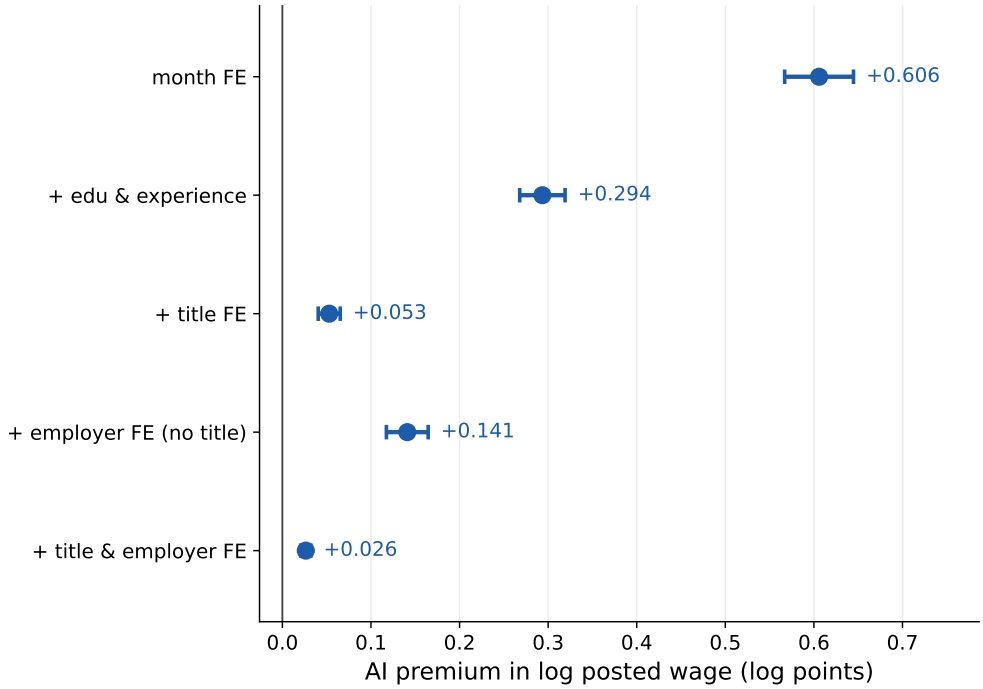
Figure D3: Posted wage event study by AI status



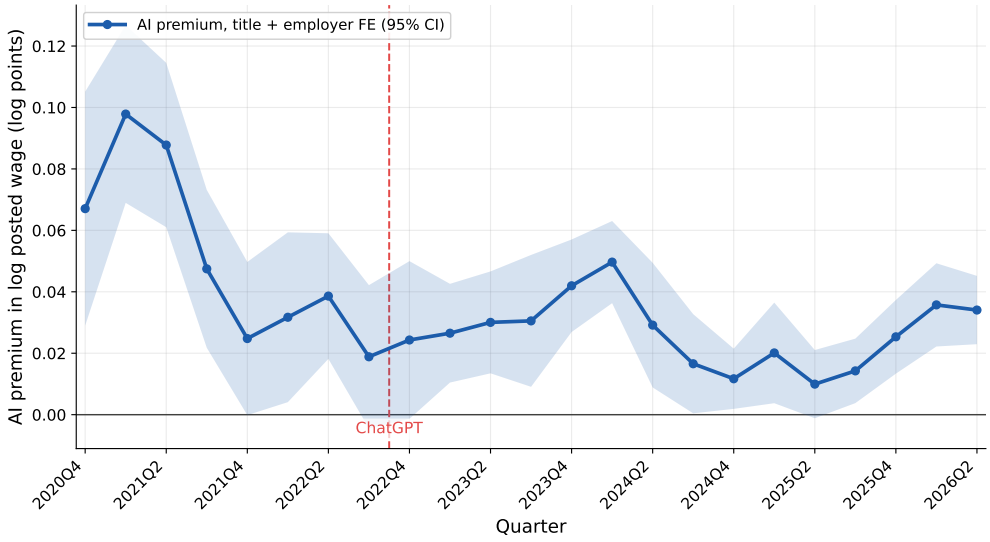
Notes: Equation 1 estimated on posted wages, with each occupation-quarter mean computed separately over ads that mention AI-related terms and ads that do not. The two series are estimated on the same occupations with the same controls and differ only in which ads enter the mean. Because ads mentioning AI are sparse in the earlier quarters, cells enter the AI series at a lower minimum ad count. Shaded bands are 95 percent confidence intervals from standard errors clustered on occupation.

Figure D4: Ad-level AI wage premium

(a) Pooled

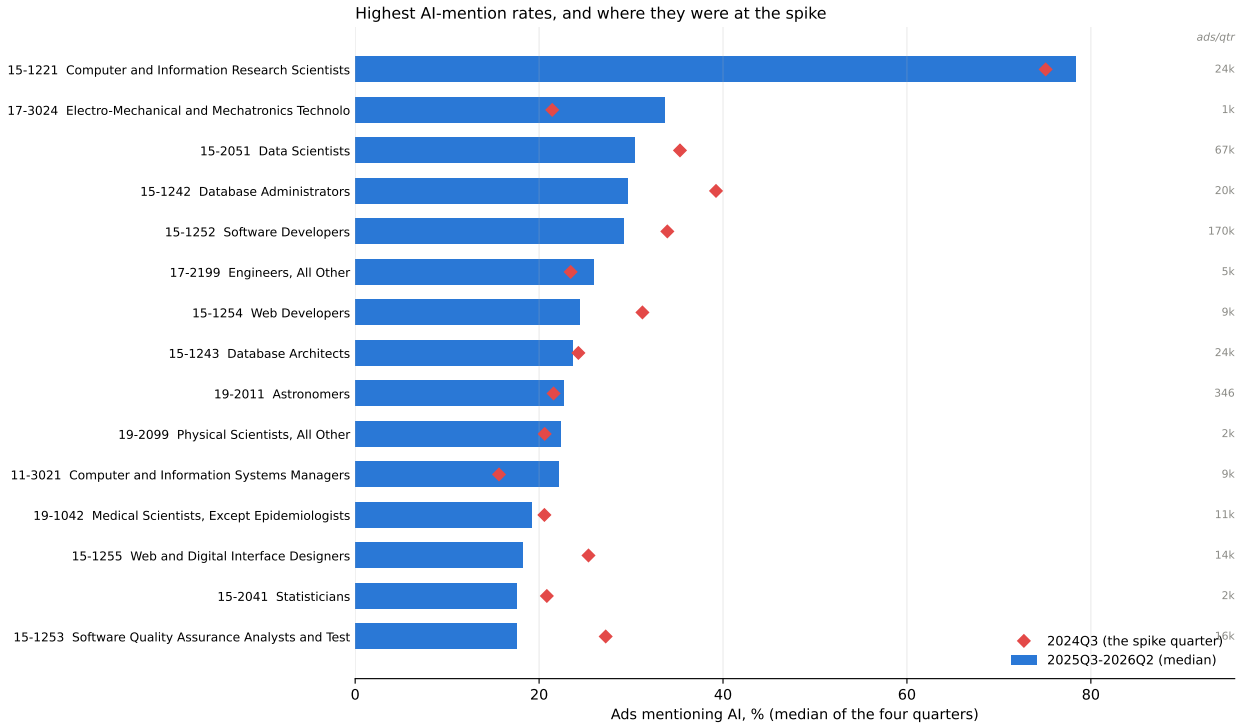


(b) Quarterly



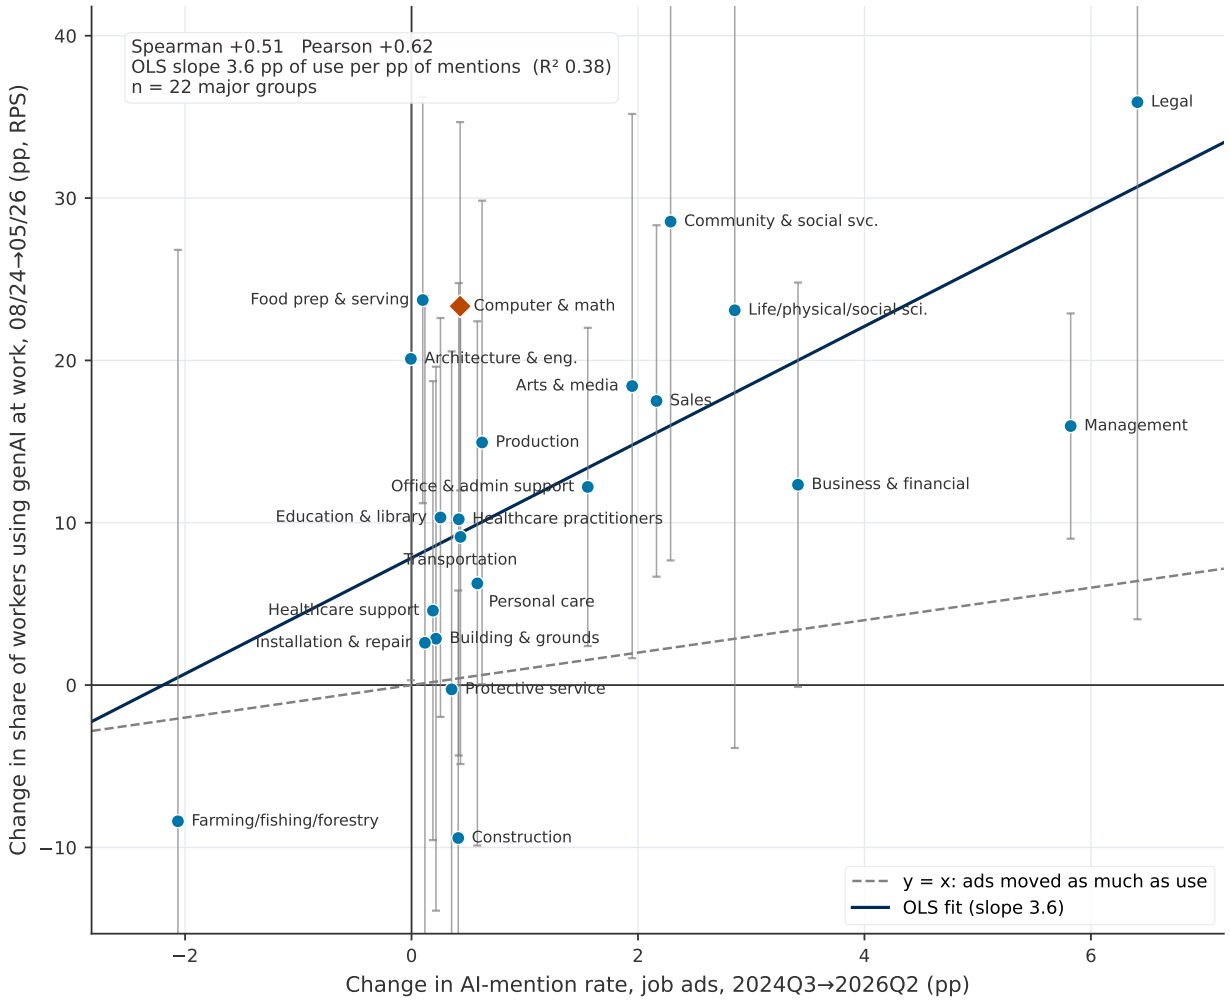
Notes: Ad-level regressions of log real posted wages on an indicator for mentioning AI-related terms, estimated on a 10 percent random sample of full-time ads that post a wage, 2020Q4–2026Q2. Panel (a) adds fixed effects cumulatively; sample sizes fall across rows as title and employer fixed effects absorb ads whose title or employer appears only once. Panel (b) reports the premium quarter by quarter under the specification with both title and employer fixed effects, so it compares ads for the same job title at the same employer. Wages are deflated by the CPI-U. Standard errors clustered on employer; 95 percent confidence intervals.

Figure D5: Top Occupations Mentioning AI-Related Terms



Notes: Occupations ranked by the median of their AI-mention shares over the last four quarters, which prevents a single quarter of heavy posting by one employer from carrying an occupation. Bars give that median and markers give the same occupation's share in 2024Q3; ads per quarter are reported beside each bar so that the size of each occupation is visible.

Figure D6: Change in AI Mentions in Job Ads vs. Change in Survey-Reported AI Use



Notes: Change from 2024Q3 to 2026Q2 in the share of employed workers reporting generative-AI use at work, from the Real-Time Population Survey of Bick et al. (2024), against the change in the share of job ads mentioning AI, by major occupation group. Ad shares count any AI reference, including generic "AI" language. Whiskers are 95 percent confidence intervals on the survey change and are clipped where they run past the plotted range; the axes have different ranges, so the 45-degree line does not appear at 45 degrees. Survey cells are small for the smaller groups, so no single group should be read on its own.