The Effect of AI on Jobs

Anahid Bauer*

Eliza Forsythe†

August 4, 2026

The rise of artificial intelligence and machine learning tools (AI/ML) has the possibility of disrupting employment on the scale of other major generational technological breakthroughs, such as electrification and the invention of the personal computer. Such disruptions have the opportunity to provide aggregate economic benefits, but often come with substantial variation in individual impacts, with some left worse off. The previous wave of technological change most affected routine jobs, resulting in negative employment and wage effects for many middle skill workers (Autor & Dorn, 2013), but also wage and employment growth for individuals with complementary skills (Bresnahan, Brynjolfsson, & Hitt, 2002; Dillender & Forsythe, in press). However, the era of generative AI threatens to replace cognitive and skilled tasks (Webb, 2019; Brynjolfsson, Mitchell, & Rock, 2018; Felten, Raj, & Seamans, 2018; Kogan, Papanikolaou, Schmidt, & Seegmiller, 2023; Eloundou, Manning, Mishkin, & Rock, 2023), exposing a new set of workers to disruption.

In this paper, we ask: how does the adoption of AI/ML technology change a firm's labor demand in terms of skill and task composition? We measure the effect of a firm's adoption on the AI-intensity of its own job ads.

Our job-ad data come from Web Data Commons (WDC), which extracts schema.org/JobPosting structured markup from the Common Crawl web corpus — the machine-readable job-posting layer that employers embed in their career pages for job-search engines. We use 11 snapshots spanning 2015–2026 (more than 8 million postings), each posting carrying the vacancy title and description, the hiring organization's name, posting date, and industry/location fields. Because schema.org provides no firm identifier, we construct one from the (cleaned) employer name. WDC coverage is limited to web pages captured by Common Crawl that carry job-posting markup, which over-represents larger, digitally mature employers — a selection we return to when interpreting external validity.

Our identification relies on a sample of companies that use ChatGPT provided by Bloomberry, which pins down when a company adopted ChatGPT Business and Enterprise by monitoring new entries and modifications to company DNS records. The advantage of relying on changes in DNS records is that this source is completely independent of job ads content. We focus on large companies of more than 500 employees with a known date of adoption.

We build on the work of Meisenbacher and Norlander (2026) to construct a list of 642 AI related words by including post-2024 GenAI vocabulary (new model names, agentic patterns), RAG/vector infrastructure terms, and AI safety/evaluation language. Given the

*Institut Mines-Telecom Business School,LITEM, Paris Saclay anahid.bauer@imt-bs.eu

†University of Illinois, Labor and Employment Relations School and Department of Economics, eforsyth@illinois.edu

fast-paced evolution of this technology, lists of AI-related words needs permanent updating to capture new vocabulary.

Our sample is restricted to ads from identifiable employers (more than 6.8M ads). For the event study we keep treated firms with a known adoption date that appear in at least two snapshots (1,787 firms), compared against a pool of more than 90K control firms also observed in at least two snapshots.

One might worry that firms' decisions to adopt generative AI and to hire workers with AI-related skills are plausibly correlated with the same unobserved firm characteristics. To separate adoption from this confounding, we exploit the timing of adoption within firms rather than cross-firm differences in whether they adopt: firm fixed effects absorb each firm's time-invariant baseline AI vocabulary and adoption propensity, and snapshot (calendar) fixed effects absorb the strong secular rise in AI language common to all firms. Identification therefore comes from the staggered timing of ChatGPT adoption — at any given snapshot, firms that have already adopted are compared with those not yet adopted.

We estimate a dynamic, staggered-adoption event study of AI-word intensity around each firm's adoption date, with event time measured in years. Our outcomes are the share of a firm's ads containing any AI term, the mean number of distinct AI keywords per ad, and the mean number of AI-word occurrences per ad. As a benchmark we first estimate a two-way fixed-effects (TWFE) event study. Under staggered adoption with heterogeneous treatment effects, however, TWFE is biased because already-treated firms enter as effective controls for later adopters ("forbidden comparisons"; Goodman-Bacon, 2021). Consistent with this, the TWFE estimates are near zero.

Our preferred specification is a matched-control difference-in-differences design that supplies clean controls from non-adopting firms. For each adopting firm we build a comparison group of five non-adopting firms — drawn from the same industry and region and closest to it in ad volume — and follow the adopting firm alongside these five over time (a stacked design). Both the adopting firms and their comparison firms are restricted to those observed in at least two snapshots. We then estimate an interacted event study with group-by-firm and group-by-period fixed effects, weighting by the number of ads per firm-snapshot and clustering standard errors by firm. As robustness checks we re-estimate the dynamics using the Callaway and Sant'Anna (2021) estimator — treated-only, with not-yet-treated firms as controls — and a matched version of the same estimator on the control group described above.

We find that after adopting generative AI, firms nearly double AI intensity in hiring across all three outcomes, with an increase of 12 percentage points in the share of a firm's ads containing any AI term, of 0.3 in the mean number of distinct AI keywords per ad, and of 0.6 in the mean number of AI-word occurrences per ad. As expected, TWFE leads to null results, since it uses already-treated firms as control observations, biasing the results downwards. The Callaway and Sant'Anna (2021) estimator shows supporting evidence of a positive effect, but the estimation is noisier given the annual gapped panel.

In ongoing work we probe the result along several dimensions. First, to confirm that the effect reflects demand for generative AI rather than a broad ramp-up in tech hiring, we re-estimate on a narrow, generative-AI-specific sub-dictionary — a 140-term subset of our keyword list restricted to large-language-model, agentic, retrieval/vector, and generative-tooling terms (dropping generic programming and data-science skills such as Python and

SQL). Second, we test whether the pattern generalizes beyond ChatGPT by extending the treatment definition to adopters of other generative-AI providers (e.g., Anthropic's Claude), which also enlarges the treated sample. Third, we test the robustness of these dynamics in Lightcast, which offers near-universe US coverage and firm identifiers; its monthly frequency lets us estimate a higher-frequency event study for the 2024 adoption cohort, and extending to later vintages would allow the same test for the 2025–2026 cohorts. Finally, we note that WDC coverage is limited to web pages captured by Common Crawl that carry job-posting markup — over-representing larger, digitally mature employers — so our estimates speak to the population of firms that advertise on the open web; the Lightcast comparison also serves to gauge this selection.

This paper makes three contributions. First, and most centrally, we propose an identification strategy for firm-level technology adoption: rather than inferring adoption from job postings themselves — which conflates the treatment with the outcome — or from surveys, we date each firm's adoption from an objective infrastructural marker, changes in its DNS records, observed at a precise date and generated entirely outside the labor market. This independence lets us exploit the staggered timing of adoption in an event-study design, and more broadly offers a template for dating firm technology adoption well beyond generative AI. Second, we provide within-firm, before-and-after evidence that adopting a general-purpose AI technology reshapes the skill content of a firm's hiring — the near-doubling of AI-related language in its ads following adoption — moving beyond the cross-sectional exposure correlations that dominate the literature. Third, we contribute a transparent, replicable measure of firm-level AI-skill demand built from public web data, including a narrow generative-AI sub-dictionary that separates genuine generative-AI demand from broader tech hiring. To our knowledge, this is among the first uses of firm-level, DNS-dated adoption as a treatment for studying technology diffusion in labor markets.

References

Autor, D. H., & Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American Economic Review*, *103*(5), 1553–1597.

Bresnahan, T. F., Brynjolfsson, E., & Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *The Quarterly Journal of Economics*, *117*(1), 339–376.

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn and what does it mean for occupations and the economy? In *Aea papers and proceedings* (Vol. 108, pp. 43–47).

Dillender, M., & Forsythe, E. (in press). Computerization of office collar jobs. *Journal of Labor Economics*.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). Gpts are gpts: An early look at the labor market impact potential of large language models. *arXiv preprint arXiv:2303.10130*.

Felten, E. W., Raj, M., & Seamans, R. (2018). A method to link advances in artificial intelligence to occupational abilities. In *Aea papers and proceedings* (Vol. 108, pp. 54–57).

Kogan, L., Papanikolaou, D., Schmidt, L. D., & Seegmiller, B. (2023). *Technology and labor displacement: Evidence from linking patents with worker-level data* (Tech. Rep.). National Bureau of Economic Research.

Meisenbacher, S., & Norlander, P. (2026). *Building a custom taxonomy of ai skills and tasks from the ground up with job postings.* Retrieved from https://arxiv.org/abs/2605.21029 doi: 10.48550/arXiv.2605.21029

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.