How Retrainable are AI-Exposed Workers?*

Benjamin Hyman† Ben Lahey‡ Karen Ni§ Laura Pilossoph¶

February 3, 2026

Abstract

As artificial intelligence (AI) capabilities advance, will workers best adapt by upskilling in AI-complementary work or reskilling into occupations less exposed to AI? To answer this question, we assemble a new historical dataset of over 1.6 million job training spells from the U.S. Workforce Investment and Opportunity Act (WIOA) linked to transitions in AI exposure before and after training. Drawing inferences from past cognitive retraining experiences and more recent AI upskilling efforts, we compare earnings outcomes for trainees with a matched sample of workers who faced similar AI exposure but received only job search assistance. We find that the average earnings returns to training are high across the AI spectrum—$1,400 to $2,600 per quarter for high- and low-exposed occupations respectively. Low-exposure occupations see stronger gains by upskilling in AI, while high-exposure occupations receive higher returns from moving toward less AI-exposed work. This suggests that a worker's optimal AI skill investment may lie in an intermediate range. Using a new AI retrainability index (AIR), we estimate that about 45%-50% of occupations are "AI-retrainable"—workers leaving these occupations are compensated with higher earnings from moving up the AI skill ladder.

JEL Codes: J08, M53, O31
Keywords: Artificial Intelligence, Active Labor Market Policies, Job Training, Labor Markets

*We thank Joe Altonji, David Card, Anders Humlum, Lisa Kahn, Fabian Lange, Michael Lee, Kyle Myers, Steve Raphael, Daniel Rock, Wilbert van der Klauuw, Till von Wachter, and seminar participants at the annual ASSA/AEA and SOLE conferences, Federal Reserve Banks of San Francisco and New York, MIT FutureTech, and UC Berkeley IRLE for helpful comments on an earlier version of the paper. We also thank Jonathan Lee for excellent research assistance. Ni gratefully acknowledges support from the Institution of Education Sciences, US Department of Education, through grant R305B150012 to Harvard University.

†University of California, Los Angeles. Email: benhyman@ucla.edu
‡New York University, Department of Economics. Email: bpl9631@nyu.edu
§Harvard Kennedy School. Email: kni@g.harvard.edu
¶Duke University, Department of Economics. Email: laura.pilossoph@duke.edu

1 Introduction

The debate over whether advances in artificial intelligence (AI) will ultimately complement or substitute labor has drawn significant attention (Autor (2024); Deming et al. (2025); Hampole et al. (2025)). However, there is a dearth of research examining the role that existing job training programs might play in helping AI-exposed workers adapt to the evolving labor market. This gap is particularly striking in light of survey evidence showing that AI-adopting firms are retraining their workforces more rapidly than they are adjusting headcounts in response to AI technologies (Abel et al., 2024).[1]

This paper estimates the extent to which the US's flagship workforce development program—the Workforce Innovation and Opportunity Act (WIOA), formerly the Workforce Investment Act (WIA)—assists workers in successfully reskilling out of occupations facing AI-related pressure and into jobs with higher AI complementarity. We assemble a novel historical dataset of over 1.6 million individual WIOA/WIA training spells from 2012 to 2023 linked to administrative earnings records spanning several quarters before and after training. The earlier years allow us to draw inferences from past cognitive training experiences in addition to the more recent program years in which AI-reskilling became more prevalent. Crucially, the data include detailed occupation codes at each stage of the training process: the worker's pre-training occupation, the targeted occupation of training, and the post-training occupation. These occupation codes allow us to link workers to measures of the *potential* AI exposure of tasks around job training spells (Eloundou et al. (2024), Brynjolfsson et al. (2018)) as well as *realized* AI exposure from task usage with Anthropic's large language model (LLM), Claude (Handa et al. (2025)). We use this framework to estimate how transitions in AI-skill space mediate the earnings returns to training, decomposing those returns into components driven by AI-upskilling versus alternative mechanisms.

We find that the average earnings returns to training are high across the AI exposure distribution, with the strongest gains among initially low-exposure workers who upskill into more AI-intensive tasks, and smaller gains among workers from highly exposed occupations who transition toward less AI-intensive work after training. In an ideal experiment, these returns would be identified by random assignment. Absent random assignment, we follow the approaches of Rothstein et al. (2022) and Andersson et al. (2024), matching each WIOA/WIA trainee to a similarly AI-exposed control worker who also sought out workforce services, but only received job search assistance. The design's credibility is strengthened by the fact that individuals are quasi-randomly assigned to WIOA

[1] In a large-scale survey of Danish workers most exposed to ChatGPT, Humlum and Vestergaard (2024) find that 43% of workers reported the need for more training as the largest barrier to adoption.

caseworkers who gatekeep entry into training based on subjective interpretations of eligibility rules. Because both training and control group units experience a similar decline in earnings prior to receiving services, the strategy also addresses the potential "Ashenfelter dip" concern that a decline in earnings prior to training could upwardly bias estimates of post-training earnings gains (as first shown by Andersson et al. (2024)). The nearest neighbor approach is also well suited to separately examine outcomes for workers displaced from occupations with varying initial AI exposures. Though not without limitations, this selection-on-observables approach shines important early light on the understudied intersection of AI and job training.

We document four main empirical findings. First, job training participants disproportionately originate from occupations that were highly exposed to AI prior to enrollment: the modal trainee is displaced from a top-quintile AI-exposure occupation despite having relatively low prior earnings. This pattern not driven by sample composition, as participants closely resemble workers in the CPS Unemployed and Displaced Worker samples. Second, the average earnings returns to training are large across the AI exposure spectrum—approximately \$1,400 to \$2,600 per quarter for workers from high- and low-exposure occupations, respectively. Third, returns differ systematically by direction of skill adjustment: workers from low-exposure occupations experience the largest gains when they upskill in AI-related tasks, while workers from high-exposure occupations benefit more from transitions toward less AI-intensive work, suggesting that the earnings-maximizing level of AI skill lies in an intermediate range. Fourth, positive returns for both groups are concentrated in the most recent years when large language models were most prevalent. However, this period coincided with historically tight labor markets, raising the possibility that training credentials carry greater signaling value when firms are hiring more aggressively and reaching deeper into the skill distribution.

Using the matched sample, we also construct an AI Retrainability Index (AIR) which ranks occupations by the share of workers who successfully retrain into higher-wage, more AI-intensive roles after leaving those occupations. Defining AI retrainability in terms of both skill and earnings helps mitigate survivorship bias: it prevents us from mistakenly labeling workers as successfully "AI retrainable" if AI-intensive firms have become more dominant in hiring. Earnings movements allow us to disentangle the source of such transitions by revealing whether workers are compensated for AI upskilling. Using the index, we estimate that about 45%-50% of occupations are "AI retrainable" as measured by its workers receiving higher pay for moving to more AI-intensive occupations—a large magnitude given a relatively low-income sample of displaced workers. The index also enables us to decompose, for any given occupation, the extent to which successful AI retraining is driven by wage gains while holding

skills constant (e.g., gains from occupational licensing barriers) versus skill transitions from AI training.

Our findings contribute to both the literature on the labor market effects of AI and the active labor market program evaluation literature. While personnel studies have found complementarities between AI adoption and worker productivity within individual firms, these studies generally do not attempt to isolate the effects of AI-upskilling incumbent workers from the outcomes of AI adoption (e.g., Brynjolfsson et al. (2025b) on ChatBots aiding less-experienced workers). They are also silent on the question of whether workers would be better off avoiding AI occupations altogether. Recent evidence also highlights the role of individual beliefs and barriers to reskilling: Delfino et al. (2026) show that jobseekers' perceptions of monetary returns and occupational fit shape their demand for retraining into high-demand jobs. To our knowledge, our paper is the first to directly estimate the effects of AI retraining on worker earnings, and to do so at a nationally-representative scale.

The implication that the labor demand environment—both evolving AI skill requirements and tight labor markets—is an important driver of positive returns also contributes to the literature on active labor market policy evaluations, especially studies focused on heterogeneity over the business cycle. In a meta-analysis of more than 200 evaluations, Card et al. (2018) provide suggestive evidence that active labor market policies are more effective in weak labor markets, consistent with the view that employers can be more selective in slack labor markets. We find evidence in favor of the opposing viewpoint that job training is less effective during recessions. Our results better align with the alternative explanation that training provides a particularly valuable signal when labor markets are tighter and firms have to recruit lower down in the skill distribution.

A secondary contribution of our work is that we document that modern AI exposure measures (e.g. Eloundou et al. (2024), Handa et al. (2025)) have a natural mapping to task content measures from the preceding automation literature (Acemoglu and Autor (2011)). Specifically, AI exposure measures align most strongly with routine cognitive task content, while correlations with non-routine cognitive tasks—especially for Claude-use based measures—are only modestly weaker. This mapping is helpful because it underscores that estimates of the returns to job training by time-invariant AI exposure measures will reflect routine cognitive skill retraining in earlier years (more aligned with automation) and increasingly non-routine cognitive skill retraining in later years (more aligned with AI). As the next phase in the long-term reallocation of cognitive work, understanding earlier evidence on cognitive skill retraining is informative for adaptation in the as yet uncertain AI era.

The paper proceeds as follows. Section 2 provides background information and descriptive statistics on the WIOA/WIA training programs and defines AI exposure measures. Section 3 details our nearest-neighbor matching identification strategy. Section 4 presents our main results. Section 5 outlines our AI retrainability (AIR) index which allows us to decompose the sources of successful AI retraining. Section 6 concludes and discusses areas for future research.

2 Data and Measurement

Determining whether individuals can be retrained for AI-related work is challenging because analysts typically lack administrative data on how AI is used by workers within firms.[2] We make progress by developing a dataset that can longitudinally track transitions of job training participants into occupations that vary in their susceptibility to AI; occupations are then categorized as AI-exposed using measures we describe below. To determine whether transitions into more AI-intensive roles reflect genuine upskilling—as opposed to repeated exposure to adverse AI-related shocks—the dataset must also include pre- and post-training earnings. For this purpose, we construct a new large-scale dataset of training participants by merging two publicly available data sources.

2.1 WIOA / WIA Participant Data

The US Department of Labor's (USDOL) Participant Individual Record Layout (PIRL) is a national performance tracking dataset that contains one observation for each unique participation spell in any WIA, or its successor WIOA, program from 2005q1 to 2023q4.[3] WIOA/WIA is a federally funded active labor market program in which qualified workers receive fully subsidized occupational job training (*Title I - "Workforce Development Activities"*) or job search assistance services (*Title III - Wagner-Peyser Act Employment Services*) among other services.[4] We restrict our analysis to participants who exited WIOA/WIA between 2012q1 and 2023q4, as occupation codes are available throughout all stages of the job training spell during this period.

[2]While recent surveys provide firm AI usage rates (Bonney et al. (2024)) and some estimates of worker-level adoption across the income distribution (Hartley et al. (2024)) and by occupation (Humlum and Vestergaard (2024)), we are aware of only one study that has linked AI use to worker administrative data (Pizzinelli et al. (2023)) though it does not include individual-level job training data.

[3]WIA data adheres to the PIRL infrastructure, while its successor WIOA uses the updated Workforce Investment Act Standardized Record Data (WIASRD) files. Files are currently updated through 2024, however not enough time has yet elapsed to observe outcomes from 2024 participants.

[4]Popular examples of training services are nursing degree programs and courses in data analysis with Microsoft Office and Python. See Andersson et al. (2024) for an involved description of available services.

Our main focus is on Title I occupational training programs which span over 1.67 million training spells from 2012 through 2023. 57% of Title I participants are enrolled in the *Title I - Adult* program which prioritizes current public service recipients (e.g. Supplemental Nutrition Assistance Program (SNAP) recipients), lower-income individuals, and disadvantaged individuals with employment barriers. 35% of Title I participants are enrolled in the *Title I - Dislocated Worker* program which serves individuals who have been laid off, received a notice of termination, or are eligible for USDOL rapid response services directed to workers impacted by mass layoffs or plant closures. Unlike Adult program trainees who can be part-time employed while training, Dislocated Worker program trainees are generally enrolled full-time in training. Of the remaining Title I participants, less than 1% are in the program for young adults aged 14-24, while 8% are co-enrolled across the above programs.

To establish a control group of similarly AI-exposed workers, we follow Rothstein et al. (2022) and use matched participants in Title III programs who receive job search assistance services through the Wagner-Peyser Act, but do not enroll in occupational retraining programs.[5] Title III programs span over 29 million participants from 2012 to 2023 providing a rich set of characteristics from which to match Title I trainees. While both Title I and Title III workers can receive job search assistance services upon initial intake at American Job Centers, individuals must be evaluated by caseworkers to enroll in Title I training services (29 U.S.C. §3174(c)(3)(A)(i)).[6] Because individuals are quasi-randomly assigned to caseworkers who may vary in how they interpret ambiguity in training eligibility rules (Appendix D), the resulting variation in training assignment reduces selection-on-unobservables concerns (Appendix E). To account for any lingering selection bias from endogenous sorting into Title I and Title III programs, we use a matched sample strategy described in detail in Section 3.

We observe individual quarterly earnings from state unemployment insurance administrative records (ES-202 forms) for three quarters before each participant enters a Title I or Title III program, and for three quarters after exiting the program. We deflate all earnings records to 2010Q1 dollars. Importantly, for a sizable subset of individuals we observe the quarterly 6-digit Standard Occupation Code (SOC) associated with each worker's primary job if employed.[7] One limitation of the PIRL dataset is that the

[5] Wagner-Peyser Act services include basic and core services—job search assistance, referrals, labor market information and resume-building workshops—but can also span intensive career services like counseling.

[6] Because the case worker is unobserved, we cannot use an examiner design to estimate effects as in Humlum et al. (2023).

[7] All 6-digit SOC codes in our data are derived from 2010 detailed O*NET (Occupational Information Network) occupational codes. SOC codes represent the first 6 digits of the 8-digit O*NET codes.

requirement to report occupation codes is left to the discretion of local workforce development boards. After restricting to participants with complete occupation codes (available for 22% of training participation spells), earnings records, and covariate information and winsorizing the top 1% of average pre-WIA/WIOA earnings, our final sample contains 116,198 total training spells, of which 95,903 spells are successfully matched to a nearest neighbor control group unit. In Section 2.3, we provide balance tests that confirm that individuals with missing occupation codes are qualitatively similar to those with observed occupation codes, consistent with arbitrarily distributed reporting requirements for occupation codes.

Finally, for each individual, the PIRL also provides a rich set of demographic variables including sex, race, ethnicity, highest level of educational attainment, disability status, and social benefit distinctions including low-income status, Temporary Assistance for Needy Families (TANF) recipient, SNAP recipient, or other public assistance recipient, which are helpful for both description and matching.

2.2 Occupation-Level AI Exposure Measures

Firm-level AI use and individual-level AI exposure are rarely observed at scale in administrative data. This measurement constraint is well recognized in the digital economics literature, and a standard response is to proxy exposure using occupation-level indices for whether an occupation's work tasks can be, or have been, accomplished by AI. To this end, we merge two leading occupational measures of AI exposure to WIOA/WIA participants' pre-training occupations to study AI-related transitions.

Eloundou et al. (2024) – potential AI exposure. The first measure, from Eloundou et al. (2024), builds on the seminal "suitability for machine learning" (SML) framework of Brynjolfsson et al. (2018). In that work, human annotators evaluated 2,069 O*NET work activities ("tasks") across 964 six-digit SOC occupations, rating each task–occupation pair on its suitability for machine learning.[8] Eloundou et al. (2024) extend this approach to large language models (LLMs), asking whether generative AI systems can substantially reduce the time required to perform specific tasks. Annotators—and ChatGPT itself—assess whether (a) an LLM alone could reduce task completion time by at least 50%, and (b) additional software built on top of the LLM could achieve a 50% time reduction.[9] A key advance in Eloundou et al. (2024) is coverage: rather than the 2,069

[8]Seven independent annotators scored each task–occupation pair on 23 Likert-scale questions capturing the likelihood that a machine could perform the task in that occupational context. A score of 5 indicates high suitability; 1 indicates low suitability.

[9]We use Eloundou et al. (2024)'s β weighting scheme, which defines AI exposure $AI^{\beta} \in [0,1] = s_a + 0.5s_b$

activities used in the original SML work, they evaluate all 18,156 granular O*NET tasks, many of which are occupation-specific. As in Brynjolfsson et al. (2018), task-level exposure is aggregated to the occupation level using O*NET task-importance weights, which are derived from surveys of job incumbents and occupational experts who report how critical each task is to performance in that occupation.

Among the unique occupations spanned by the WIOA/WIA dataset, the mean and median AI exposure scores (prior to program participation) are 0.339 and 0.342 with a standard deviation of 0.223. The distribution of AI exposure scores across all 6-digit occupations in the Eloundou et al. (2024) dataset (not conditioning on training participation) is very similar, having a mean of 0.326, a median of 0.327, and a standard deviation of 0.220. Consistent with Brynjolfsson et al. (2025a), highly AI-exposed customer service representatives are among the most prominent participants in job training, while the lowest exposure quintile is dominated by manual occupations that are comparatively insulated from AI. Table 1 illustrates this pattern.

Anthropic Economic Index (2025) – realized AI exposure. A limitation of the Eloundou et al. (2024) exposure measure is that it may not reflect how AI is being used in practice since it is a measure of what tasks LLMs could *potentially* complete. To fill this gap, we thus also draw on *realized* usage data from Anthropic's Economic Index (Handa et al. (2025)). This index estimates task-level LLM usage using a random sample of one million Claude conversations in late 2024 and early 2025. Its key advantage is that it reflects observed behavior rather than hypothetical task feasibility.[10] To map task-level Claude usage to occupations, we follow Anthropic's preferred allocation procedure, dividing total usage on a task by the number of occupations that perform that task. In practice, most tasks are occupation-specific, so this adjustment affects only a minority of cases. To ensure comparability with Eloundou et al. (2024), we further weight tasks by their O*NET task-importance scores when aggregating to the occupation level. Analogous to Table 1, Appendix Table A.7 reports the three most prominent occupations among training participants within each quintile of Claude-based AI exposure.

AI Exposure versus Automation

In this paper, occupational measures of AI exposure reflect the degree to which AI technology *could* perform the same tasks previously done by a human. AI exposure should

where $s_a, s_b \in [0, 1]$ denote the shares of an occupation's tasks for which conditions (a) and (b) hold. This measure captures the share of an occupation's tasks that are LLM-exposed, and is validated by its positive correlation with average occupational earnings.

[10] A limitation is that Anthropic's user base may be disproportionately concentrated in certain industries, such as software development, which may not be fully representative of the national workforce.

Table 1: Top Occupations Prior to Training by AI Exposure

Rank	6-Digit SOC	Occupation	No. Trainees
AI Exposure Quintile 5 (High Exposure)			
1	434051	Customer service representatives	3,926
2	439061	Office clerks, general	1,377
3	436011	Executive secretaries and administrative assistants	785
AI Exposure Quintile 4			
1	533033	Light truck drivers	1,294
2	435081	Stockers and order fillers	1,283
3	119199	Managers, all other	1,206
AI Exposure Quintile 3			
1	412011	Cashiers	1,943
2	533032	Heavy and tractor-trailer truck drivers	1,914
3	412031	Retail salespersons	1,360
AI Exposure Quintile 2			
1	311014	Nursing Assistants	4,981
2	311011	Home health aides	1,138
3	353031	Waiters and waitresses	982
AI Exposure Quintile 1 (Low Exposure)			
1	519199	Production workers, all other	2,564
2	537062	Laborers and freight stock, and material movers, hand	1,923
3	519198	Helpers - production workers	1,549

Notes: This table shows the top three most prominent occupations in different AI exposure quintiles prior to entering WIOA/WIA training participation. Counts of number of training participants are shown in the final column. Table uses AI exposure measures from Eloundou et al. (2024). See text for details.

thus be thought of as the potential rather than realized substitution of AI for human tasks, though recent work has shown a high correlation between our chosen AI exposure measures and real task usage using Microsoft Bing Copilot (Tomlinson et al. (2025)). Because the PIRL data's time span from 2012 - 2023 includes earlier precursors to modern AI, some of what the above AI exposure measures capture is automation or computerization in earlier years, which are informative for the AI era. One novel contribution of our paper is that we show that two prominent exposure measures used in this literature ($AI^{\beta}_{Eloundou}$ and AI_{Claude}) correlate with foundational measures used in the Acemoglu and Autor (2011) task framework.

In Appendix Figure A.1, we illustrate this link with binscatter plots which reveal that both AI exposure measures align strongly with routine cognitive task content, while correlations with non-routine cognitive tasks—especially for Claude-based measures—are only modestly weaker. That is, we formally document that AI tasks are cognitive tasks, with non-routine cognitive tasks having a lower ceiling and realized task use with LLMs lagging potential use. These facts are useful because they underscore that estimates of the returns to job training by time-invariant AI exposure measures will reflect largely routine cognitive skill retraining in earlier years (more aligned with previous experiences with automation) and may potentially move toward non-routine cognitive skill retraining in later years (more aligned with AI).

2.3 Descriptive Statistics and Sample Restrictions

Throughout the paper, we split the sample of training participants into a "high" and "low" AI exposure group based on whether their pre-enrollment occupation was above or below the 55th percentile of the 6-digit occupation AI exposure score distribution. As we discuss in Section 3, this cutoff is determined empirically based on the stability of estimates across potential cutoffs, however all estimates are robust to shifting the cutoff in either direction.[11] Appendix Table A.1 presents means and standard deviations of key worker characteristics at the time of training enrollment for these two groups.

Among the 95,903 pooled training spells in our final matched sample, observation counts are relatively evenly distributed between high- and low-AI exposure workers prior to training, pointing to broad representation of occupations in the WIOA/WIA data. The average worker pooled across both samples made around $40,000 in annualized 2010 earnings prior to training participation. In Appendix Figure A.2, we also show that these

[11] We assign each training spell a single prior- and post-training occupation based on the SOC code associated with the plurality of earnings in the three quarters before or after training. For occupation codes that are self-reported at more aggregate levels, we use the average of all 6-digit codes within that occupation group.

workers had pre-separation AI exposure distributions that are almost identical to the distributions among the Current Population Survey (CPS) of unemployed workers when comparing occupations prior to separation in repeat cross sections over time. This suggests that WIOA/WIA training participants are likely representative of the national population of unemployed workers.[12] Appendix Table A.1 also reveals that the high AI exposure group has a larger female share (consistent with Bollinger and Troske (2025) who document greater success among females in coding-focused training programs), is more educated, and has slightly higher earnings prior to participation. The high AI exposure group is also more skewed toward participants in the dislocated worker program versus the adult (disadvantaged) worker program.

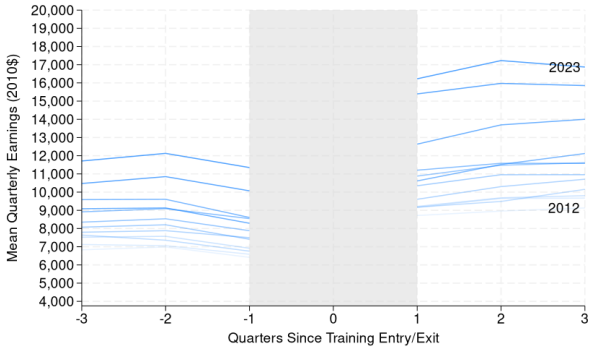
Sample Restrictions. In our summary statistics table (Appendix Table A.1) and throughout the analysis, we restrict to workers who have: (1) positive earnings in all six quarters of observation, (2) non-missing occupation codes in all six quarters of observation, (3) non-missing covariates used for matching. The first restriction intentionally limits the sample to reemployed workers, allowing us to measure earnings responses to changes in the AI content of occupations insulated from any confounding extensive margin reemployment effects. Relaxing this restriction results in near-identical earnings returns to WIOA/WIA (see Appendix Figure A.4). The second restriction ensures we can measure flows in AI skill space, however, results in dropping a significant share of training participants. This is because the reporting of occupation codes varies by state and may be missing "due to language in the WIOA statute and a variety of practical/quality reasons" or if the "occupation code is not a part of the UI wage record in most states" (correspondence with USDOL Employment and Training Administration, November 2024). In Appendix Table A.2, we show that compositional differences between workers in states with and without recorded occupation codes are mostly negligible. Ultimately, our matching estimator will compare training participants with nearest-neighbor matched workers within occupation bins while matching on state location.

Earnings Selection Patterns. Before turning to our matching design, Figure 1 presents selection patterns in earnings around training participation spells by calendar year of program exit subject to the above sampling restrictions. Panels (A) and (B) show real quarterly earnings for the three quarters before training program entry and three quarters after training program exit, separately by low and high AI exposure in the worker's occupation prior to participation based on the SML occupational measure.

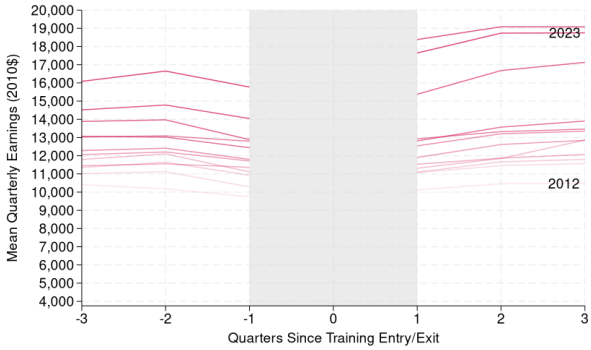
[12]Occupations spanned by our full set of WIOA/WIA training participants account for around 93% of the CPS labor force on a monthly basis.

Figure 1: Earnings Patterns by AI Exposure Prior to Job Training

(A) **Low** AI Exposure in Prior Occ.



(B) **High** AI Exposure in Prior Occ.



Notes: Training participants are classified as "high" AI exposure if their 6-digit SOC exposure score is above the 55th percentile of the Eloundou et al. (2024) 6-digit AI exposure score distribution, and "low" if below the 55th percentile. Heavier lines correspond to more recent years. The shaded gray region emphasizes that quarterly earnings are observed in the quarters prior to entry, and after exit, but not during training. All earnings are in real 2010 dollars.

The primary selection concern in the active labor market literature is the "Ashenfelter dip": earnings declines prior to training that can mechanically upward bias comparisons with post-training earnings (Ashenfelter, 1978). The event-study patterns show a modest dip in earnings immediately before training, followed by weakly higher earnings in most years and substantially larger post-training differences in more recent cohorts (shown with heavier lines), when AI technologies were most prevalent.[13] Participants from highly AI-exposed occupations begin with higher baseline earnings but display the same overall time pattern—larger post-training earnings differences in later years. The data do indicate slightly lower post-training earnings among highly exposed groups, but this difference is small relative to the dominant time pattern.

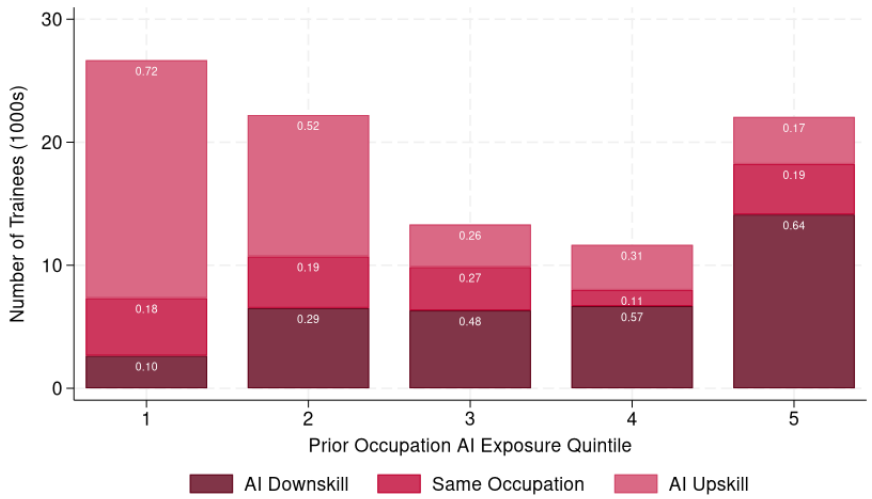
Our main strategy to handle the Ashenfelter dip is to compare training paricipants with nearest-neighbor matched program entrants that only receive job search assistance, potentially for idiosyncratic reasons related to caseworker assignments (see Appendix D). Indeed, when we plot the same patterns for matched workers who receive only job search assistance, they exhibit an almost identical pre-enrollment earnings dip (see Appendix Figure A.6). This indicates that the dip reflects shared job instability upon service entry rather than differential selection into training. These patterns suggest that Ashenfelter-type selection is unlikely to be the primary driver of the observed post-training earnings returns we find using the nearest-neighbor matching strategy described in Section 3 below.

Job Flows Across AI Skill Space. Prior to jointly analyzing the earnings returns to

[13]The stronger patterns in later years may also reflect tighter labor markets, in which training credentials carry greater signaling value when firms must reach further down the skill distribution.

training coupled with transitions in AI skills across occupations, Figure 2 first describes the flows of job trainees in AI skill space. On the x-axis, we split trainees in our sample into five quintiles of AI exposure prior to training participation based on the Eloundou et al. (2024) AI exposure of occupations. We then classify whether workers moved to occupations of lower (AI downskill), higher (AI upskill) or the same AI exposure.

Figure 2: Transitions in AI Skill Space Before and After Training



Notes: AI exposure is based on Eloundou et al. (2024) occupation-level AI exposure scores. The flows characterized in the figure correspond to the main analysis sample of the paper.

Of note, there are large shares of both highly exposed and low exposure workers among job training participants.[14] A substantial share of workers displaced from from top 20th percentile (i.e. top quintile) occupations in AI exposure prior to training enrollment. This underscores that in the context of workers interfacing with public job training programs, WIOA/WIA trainees are highly exposed to AI forces despite their lower income and education levels. Consistent with this finding, prominent occupations in the top exposure quintile include trainees who were formerly customer service representatives, cashiers, and office clerks (as shown in Table 1).

Further decomposition of the bars into the AI skills associated with destination occupations reveals that initially low AI exposed workers disproportionately upskill to higher AI content occupations, while high AI exposed workers move down in AI exposure.[15] Although part of this pattern reflects mechanical floors and ceilings of the index, it also mirrors the structure of task demand: AI appears to complement mid-level

[14] Hartley et al. (2024) find a similar bimodal pattern among users of generative AI tools.

[15] In Appendix Figure A.3, we show that while the share of workers flowing to higher AI occupations has remained stable over time, those making the transition have received growing earnings returns over time.

cognitive work in ways that reshape occupational sorting. Because AI exposure is strongly correlated with earnings, these flows imply movement toward the middle of the skill and wage distribution, consistent with AI having a leveling effect on inequality as discussed in Autor (2024). Overall, these flows point to substantial mobility across occupations after training. In subsequent sections, we disentangle the extent to which this reflects successful upskilling or instead alternative explanations.

3 Empirical Methodology

While positive earnings selection patterns and a high degree of AI upskilling are consistent with successful AI retraining, WIOA/WIA workers may be positively selected, either by virtue of the Ashenfelter dip which depresses earnings prior to training participation, or due to selection into occupational transitions. To address these concerns, we follow Rothstein et al. (2022) who use a matching design to compare WIOA/WIA trainees with Wagner-Peyser job search assistance recipients that are observationally similar, but do not enroll in occupational retraining programs. The key assumption is that, in the absence of training, both earnings and occupational trajectories for trainees would have evolved similarly to those of the matched job-search only group.

The credibility of this design is strengthened by the fact that, although all participants can receive job search assistance, entry into training is mediated by caseworkers, introducing quasi-random element to the assignment into training (see details in Appendix D and Appendix E). To control for any lingering selection concerns, our control group of nearest-neighbor matched workers who only receive job search assistance further restricts the donor pool of control workers to begin in occupations with similar AI exposure and for each nearest-neighbor pair to exit their programs in the same calendar year—the latter of which controls for the labor demand environment.

Because our goal is to attain heterogeneous estimates across AI exposure subgroups, we opt for a simple nearest-neighbor matching strategy that matches each training participant in our sample to a unique Wagner-Peyser participant, and then pools estimates across nearest-neighbor pairs within AI exposure subgroups. We show in a potential outcomes framework in Appendix F that heterogeneous comparisons across groups help difference out any remaining selection bias from the matching procedure in any given group.

We use the semiparametic matching estimator described in Abadie and Imbens (2002), where we desire a treatment effect for each treated trainee i, $\theta_i = Y_i(1)–Y_i(0)$, but only one potential outcome is observable for each unit. For implementation, we use the algorithm discussed in Abadie et al. (2004) with a Mahalanobis distance statistic between potential

units which normalizes covariates to be on similar scales and produces a single distance score between any two match candidates (see Appendix G for details). The estimator has the useful property that violators of a geometric version of the "covariate overlap" assumption are dropped, resulting in treated units only being used if the covariates share a sufficiently common covariate support. That is, large distance candidates are dropped. In our preferred sample, we add more stringent matching requirements to allow us to credibly examine many subgroups (details below). This results in dropping about 17% of the initial sample due to overlap, though less restrictive matching requirements for the overall sample would have reduced the share dropped to just 1%. This high match success rate arises due to the abundance of potential control units available in the Wagner-Peyser program.

We match each trainee to a Wagner-Peyser nearest-neighbor by choosing pairs that minimize pairwise distance calculated from a set of "fuzzy" match covariates with the requirement that a matched pair belong to the same "exact" pre-program occupation group, quintile of AI exposure, and vigintile of pre-program earnings, and exit WIA/WIOA in the same calendar year. These include:

- *Exact Match Covariates*: calendar year of program exit (controls for demand environment), AI exposure quintile in prior occupation (ensures nearest-neighbor pairs are comparable *within* similar AI exposure bins), SOC major occupation group (first two digits of 6-digit code) in prior occupation, and vigintile of average quarterly earnings before entry

- *Fuzzy Match Covariates*: sex, age at program enrollment, race and ethnicity dummies, adult vs. dislocated program indicator, limited-English flag, low-income flag, Temporary Assistance for Needy Families (TANF) recipient, veteran status, Supplemental Security Income and Social Security Disability Insurance (SSI/SSDI) recipient status, highest education attained at time of program enrollment, state centroid longitude/latitude as proxy for geographic distance of program location, quarterly pre-WIA/WIOA earnings, and average pre-WIA/WIOA earnings.

Both sets include measures of prior earnings as these have been shown in the literature to improve matching performance (Lechner et al. (2011)), and in our case, allow us to credibly estimate separate effects across the distribution of prior AI exposure. In Appendix Figure A.5, we demonstrate that both our main results and the earnings similarity between trainees and their nearest-neighbors are robust to matching without any earnings-related covariates.

After each training participation spell is matched to a nearest-neighbor job search assistance spell, we estimate balance tests across predetermined covariates in our dataset.

Appendix Table A.4 and Appendix Table A.3 present the results from these balance tests for high and low prior AI exposure samples respectively. While there are some imbalances (as would be expected by chance when estimating 44 t-tests), differences in mean values for treated and control groups are, for the most part, economically trivial or quantitatively small. The exception to this is that two variables—*employment at participation* and *total quarters enrolled*—are mechanically different. That is, trainees are more likely to be partially employed at time of participation versus their matched pairs who only receive job search assistance services, and thus remain enrolled in the program for about two fewer quarters. This longer training duration rules out the possibility that positive earnings returns from training programs are driven by quicker reemployment, as longer program durations would lead us to understate rather than overstate effects.

Mediation Analysis. The paper contains two related designs, expressed in Appendix B. The main design is a simple nearest-neighbor matching estimator based on prior AI exposure. These estimates tell us whether there is scope for AI-exposed job losers to successfully retrain and how they fare relative to non-AI exposed job losers. The second is relevant because it tells us how the choice of training program, and ultimately the destination occupation's AI exposure, matters for earnings returns. This is used when considering the share of earnings returns that can be ascribed to AI upskilling. The key insight is that a worker now has a third choice column—the target occupation AI content—which can determine outcomes. The second diagram also highlights that the nearest-neighbor assumption is more demanding in this mediation setting, because the counterfactual occupation choice absent training is unobserved and must be proxied by the control unit's realized occupation. The identifying assumption is therefore that, had trainees not exercised their training-induced choice, they would have sorted into occupations similar to those of their matched Wagner–Peyser counterparts.

4 Main Results

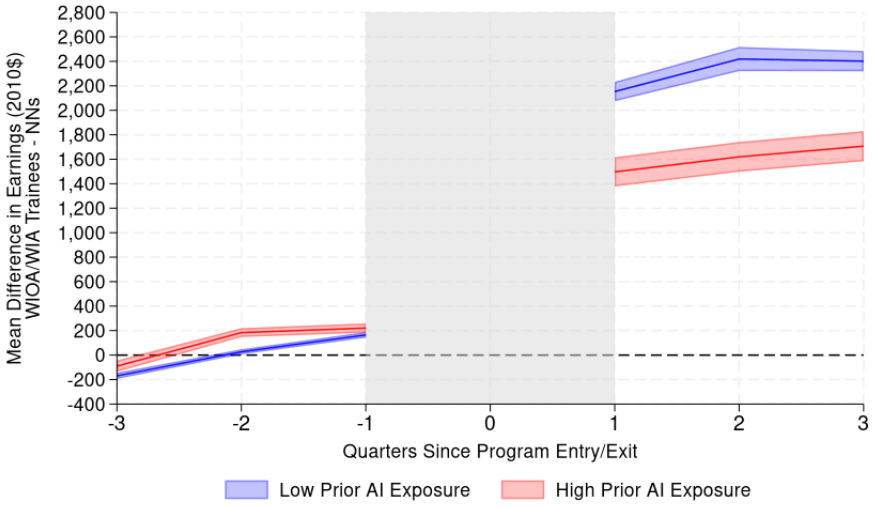
4.1 Returns to Training by AI Exposure.

In Figure 3, we graphically present the main empirical findings of the paper. Each series shows estimates of quarterly differences in mean earnings between WIOA/WIA training participants and their nearest-neighbor matched Wagner-Peyser job search assistance recipients separately by each AI exposure subgroup. Panel (A) shows estimates for high (red) and low (blue) AI exposure prior to program enrollment, separately estimated in each quarter with heteroskedasticity-robust t-tests along with 95% confidence intervals. The

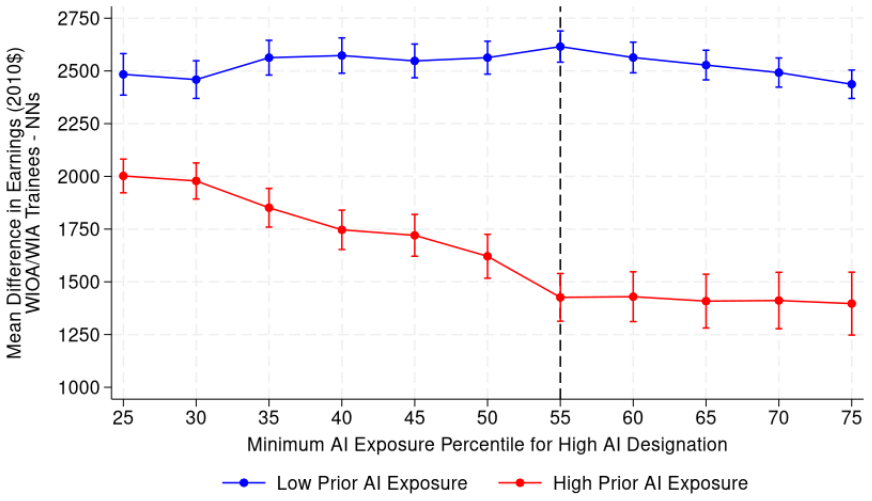
second diagram also highlights that the nearest-neighbor assumption is more demanding in this mediation setting, because the counterfactual occupation choice absent training is unobserved and must be proxied by the control unit's realized occupation. The identifying assumption is therefore that, had trainees not exercised their training-induced choice, they would have sorted into occupations similar to those of their matched Wagner–Peyser counterparts.

Figure 3: NN-Matched Earnings Returns to Job Training by Prior AI Exposure

(A) Preferred Sample Cutoff for "High" AI Exposure (55th percentile)



(B) Robustness of Pooled Earnings Returns to Choice of Cutoff



Notes: AI exposure is based on Eloundou et al. (2024) occupation-level AI exposure scores. Each series in Panel A represents the earnings return relative to nearest-neighbor matched twins. Panel B shows mean estimates from pooling together quarterly earnings across the three quarters after program exit.

Panel A shows that prior earnings differences are well balanced in both sub-groups,

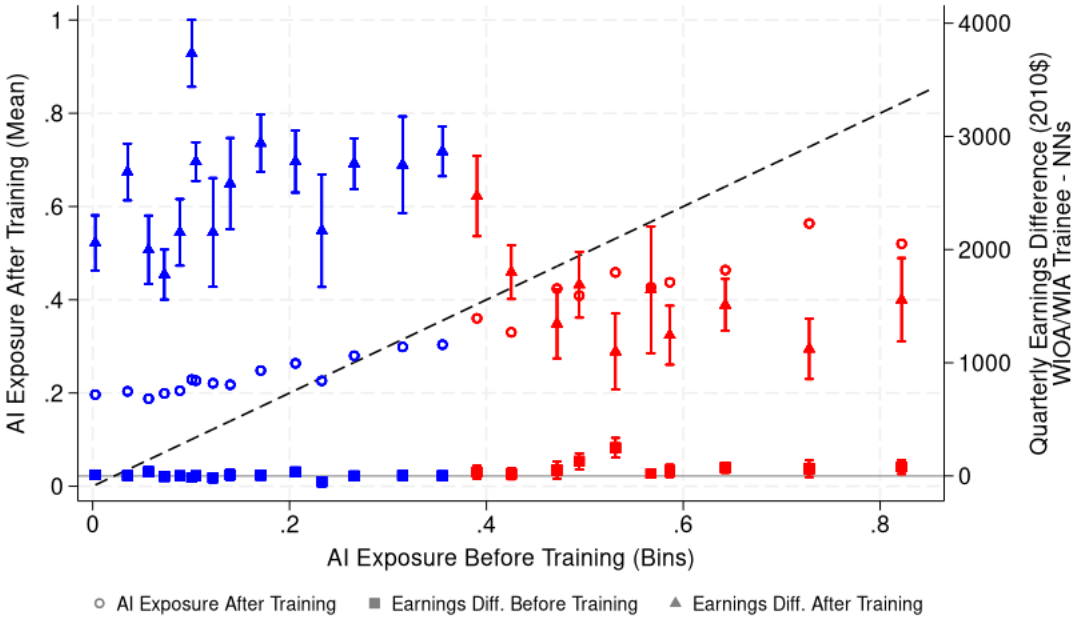
even though we are only coarsely matching on vigintiles of prior earnings. This mitigates concerns about selection on unobserved heterogeneity, and is consistent with caseworkers adding an element of quasi-random gatekeeping into training. Estimated earnings returns for high initial AI exposure training participants are about 33% lower relative to low AI exposure participants in the three quarters after program exit. However, overall training returns in both groups remain quite high in quantitative terms. In the three quarters after training exit, high AI exposure workers have quarterly earnings that are $1,313-1,539 higher relative to their matched control group pairs, or on average $5,704 annually on a base of roughly $49,000 prior to participation (see Appendix Table A.1 for baseline real wages). Among high AI exposure workers, this translates to those who train having about 12% higher short-run earnings relative to similar workers who only receive job search services. This is surprisingly large, especially considering earlier evidence suggesting that training returns were driven by the most recent years. Appendix Table A.5 shows pooled regression estimates associated with these earnings returns before and after participation. Panel B further shows that the patterns of both large effects and stronger effects for low initially AI-exposed workers is robust across all cutoffs defining exposure in the inter-quartile range.

4.2 The Role of AI Reskilling

Training participants in our main estimation strategy are not restricted to remain in their prior occupation after training. They may pursue general skill training that enables transitions away from AI-intensive occupations that face higher displacement risk, or more occupation-specific training that deepens human capital in ways that are complementary to AI rather than substitutable with it. In this section we consider how AI skill movements mediate the earnings returns from training.

Figure 4 overlays the earnings returns across the distribution of pre-training AI exposure (triangles) with estimates of workers' AI skill transitions (circles). The figure also provides pooled placebo estimates of the return to training prior to entering any workforce program (squares) to underscore that we continue to see no evidence of selection even in these more granular bins of AI exposure. Each bin plots post-training AI exposure relative to pre-training exposure, allowing us to distinguish workers who AI-upskill on average—those whose post-training AI skill content lies above the 45-degree line—from those who AI-downskill on average, whose post-training AI content falls below it. The figure therefore links earnings gains not only to where workers start in the AI exposure distribution, but to the direction of their movement in AI skill space following training.

Figure 4: Earnings Returns to Job Training Mediated by AI Skill Transitions



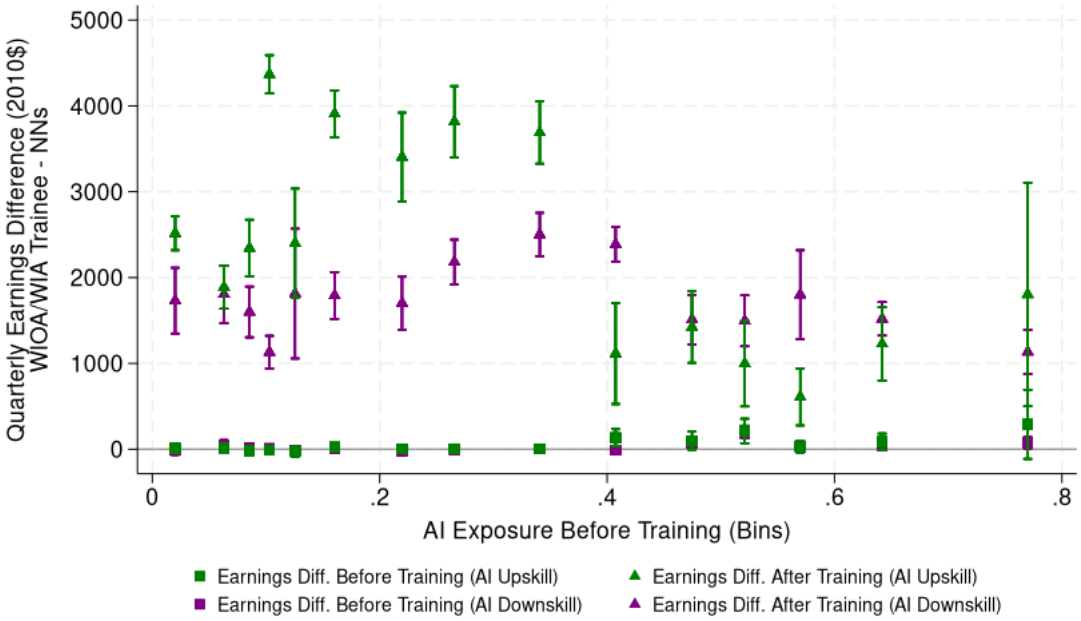
Notes: Figure overlays pooled estimates for the returns to training across the AI exposure spectrum (triangles), placebo estimates for pooled estimates prior to entry into training/job search services for the control group (squares), and average movements in AI skill before and after training (circles). Circles above the 45 degree line represent AI upskilling, while below the line represet AI downskilling. Red and blue correspond to our prior preferred high and low exposure distinction. AI exposure before training is split into 25 bins of roughly equal size, and uses the measure from Eloundou et al. (2024).

We find strong asymmetries in returns by initial AI exposure and the direction of skill adjustment. Workers originating in low-AI-exposure occupations experience the largest earnings gains when they move into more AI-intensive work, consistent with positive returns to AI-related upskilling from a low baseline. In contrast, workers from highly AI-exposed occupations earn higher returns when they transition toward less AI-intensive jobs. Together, these patterns suggest that the earnings-maximizing level of AI skill lies in an intermediate range rather than at the extremes of the AI exposure distribution. This pattern is consistent with diminishing marginal returns to AI-task intensity in workers' skill portfolios: AI skills are valuable at the margin for workers starting from low exposure, but for those already concentrated in highly AI-intensive roles, further specialization appears less rewarded than diversification toward tasks that are less directly exposed to AI.

Another way to cut the data is to separately examine those who upskilled versus those who downskilled, across the entire distribution of prior AI exposure, rather than focusing on whether the overall mean of trainees moved up in AI skill. The plot in Figure 5 does just this, and underscores that while those who AI upskill show higher returns when the starting point is low in AI exposure, this is not the case at high exposure levels. Even those who

manage to weakly move up in AI skill space do not show any further returns beyond those who AI downskill, among those highly exposed in the first place.

Figure 5: Earnings Returns to Job Training by AI Upskilling vs. Downskilling



Notes: Figure overlays separate series for training participants who upskill in AI space and participants who downskill in AI space across the initial exposure spectrum. Pooled estimates are shown for the returns to training across the AI exposure spectrum (triangles), and placebo estimates for pooled estimates prior to entry into training/job search services for the control group (squares). AI exposure prior to training is split into 15 bins of roughly equal size, and uses the measure from Eloundou et al. (2024).

Taken together, the results suggest that the payoff to AI skills is highly state-dependent: training succeeds not by universally pushing workers toward more AI, but by moving them toward the right position in the AI skill distribution given where they start.

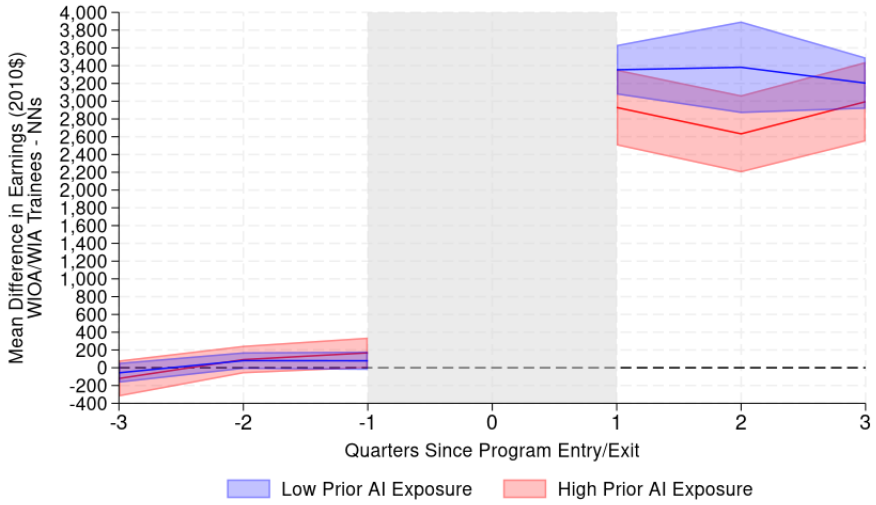
4.3 Returns to Training in the AI Era (2022+)

We have thus far considered effects for job training participants who enroll in job training programs pooled from 2012 to 2023, with the last trainee with fully populated post-training earnings entering their training programs in 2023Q4. This begs the question of how much relevance do historical returns to cognitive training have in the modern AI era?

To make progress on this question, we first re-estimate the returns to training restricted to training participants enrolled in programs after 2022q1 (such that their earnings records span the public GPT-3.5 release in late 2022) and 2023q4. Figure 6 shows that the returns to training are much higher in the post-GPT era, with a similar pattern of slightly muted

effects for those that are initially highly exposed.

Figure 6: NN-Matched Earnings Returns to Job Training (2022+)



Notes: AI exposure is based on Eloundou et al. (2024) occupation-level AI exposure scores. Each series in Panel A represents the earnings return relative to nearest-neighbor matched twins. Panel B shows mean estimates from pooling together quarterly earnings across the three quarters after program exit.

While this pattern is consistent with rising returns to training in the AI era, the same period was characterized by exceptionally tight labor markets. As a result, these estimates cannot yet disentangle AI-related adjustment from a more general mechanism whereby training has higher signaling value when firms are forced to hire further down the skill ladder.[16] While it is still helpful to draw lessons from past cognitive training experiences for the current AI environment, understanding what kinds of skills training programs increasingly emphasize remains central to interpreting these patterns.

While we do not observe the specific programs in which training participants enroll, we can characterize the evolution of AI content in workforce development programs using textual program descriptions from the WIOA eligible training provider lists (ETPLs)[17]. This can shed light on how prevalent technological training became over this period. These WIOA ETPLs contain detailed curricular and aggregate WIOA performance information on approved training programs that entered between 2014 and 2024. We use an NLP to identify common AI- and general technology-related keywords within a subsample of WIOA ETPL program descriptions, and apply those lists of keywords to tag training programs as "AI"

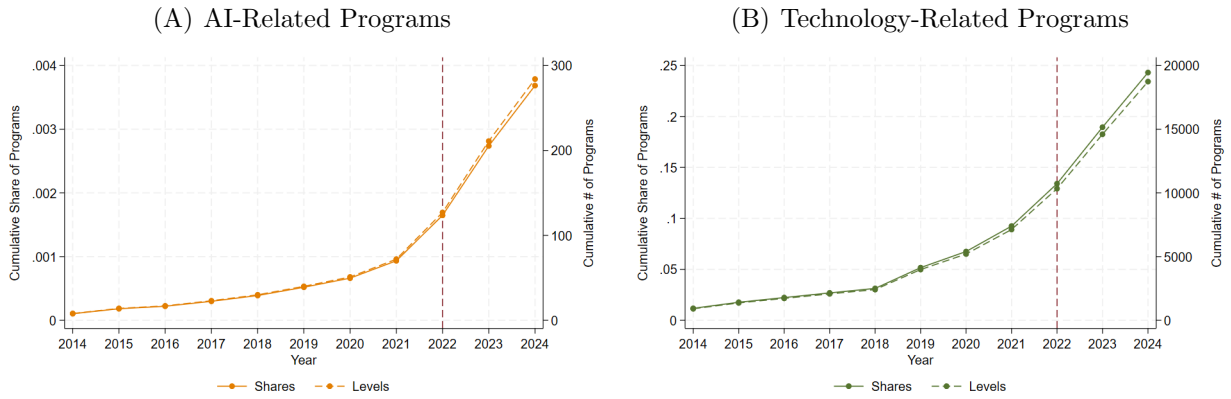
[16] In complementary work, Hyman et al. (2025) use a mismatch model following the work of Şahin et al. (2014) to quantify the degree to which the recent effects emanate from a strong labor market, or from improved alignment between the demand and supply for specific skills (especially AI skills) achieved via job training programs.

[17] The WIOA ETPL data is provided by the U.S. Department of Labor on TrainingProviderResults.gov.

related, "Technology" related, or "other" across the full dataset of WIOA training programs (see Appendix C for details).

In Figure 7, we show the growth in AI-specific and general Technology related training programs in terms of their cumulative levels and shares. Panel B shows that the share of general technology trainings grew to account for nearly 25% of all WIOA-eligible training programs by end of 2024, representing a significant share of the overall job training landscape. Although AI-specific training content still remains a very small share of all job training programs (well below the number of general technology-related programs), these programs have been growing at a rapid rate since the 2022 GPT boom. At the end of 2021, there were only about 70 total WIOA programs that contained AI-content, yet just 3 years later, the total number of AI programs nearly quadrupled. These growth trends are not just evident in the number of AI and technology training programs in the post-2022 period, but are also reflected in the number of WIOA training *participants* who enrolled in such programs (which we observe on an aggregate program basis, but not by individuals). Appendix Figure A.7 shows similar trend breaks at 2022 in AI-related programs when we weight these figures by the number of WIOA trainee participants.

Figure 7: Cumulative Growth in U.S. AI- and Technology-Related Job Training Programs

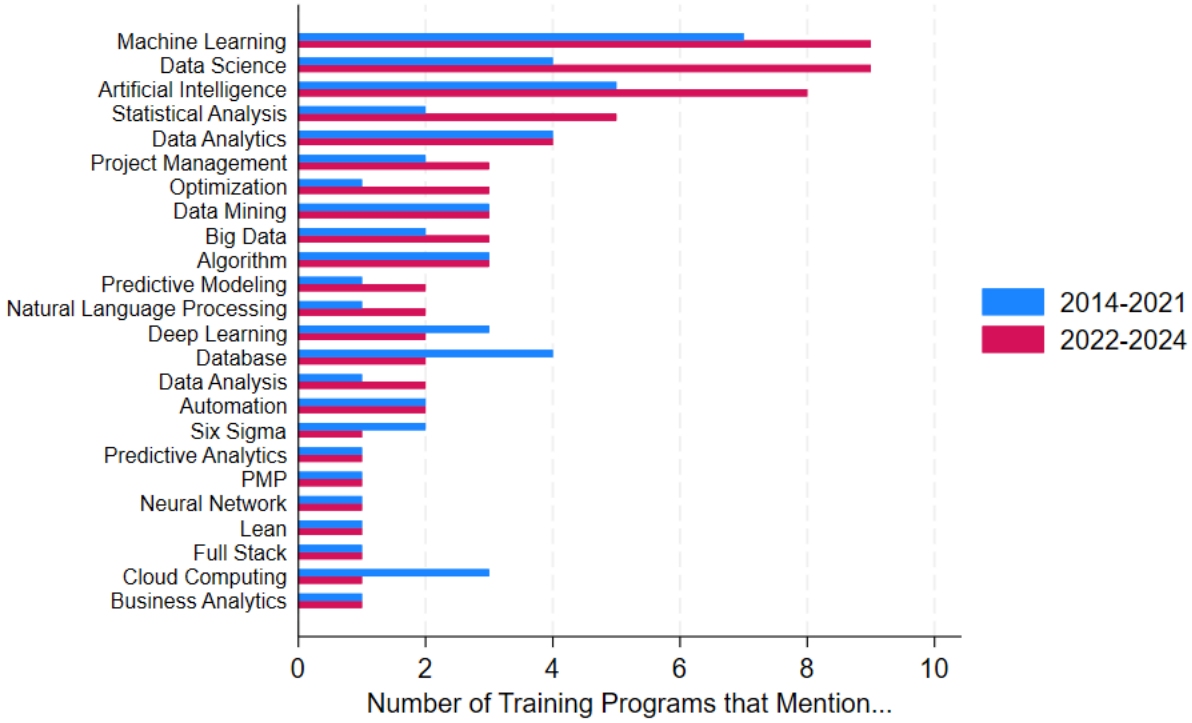


Notes: Panel A shows AI-related programs (e.g. keywords: *Artificial Intelligence, Natural Language Processing, Machine Learning*), and Panel B shows non-AI general Technology-related programs (e.g. keywords: *Programming, Web Development, Python*) drawn from WIOA ETPL lists over time. Shares are calculated as the cumulative number of AI or Technology programs introduced over the cumulative number of all WIOA ETPL programs within a given year. See Appendix C for details.

Beyond the general growth of AI-specific training programs in the post-2022 era, we also seek to understand the evolution of AI content included within such training programs. To do this, we analyze the textual descriptions of program curricula and targeted competencies within the WIOA ETPL database to draw insights on the trends in AI training content over time. Figure 8 shows a tabulation of the number of WIOA training programs containing

different AI-related keywords across the pre-2022 period (2014Q1-2021Q4) and the post-2022 period (2022Q1-2024Q4). Across the board, there are more mentions of core AI terms such as "machine learning" and "artificial intelligence" within the post-2022 period. These results highlight a growing supply in AI content by training providers as well as a growing interest in AI-training take up by WIOA participants, sharply coinciding with the advent of GPT-3.5's public release in 2022.

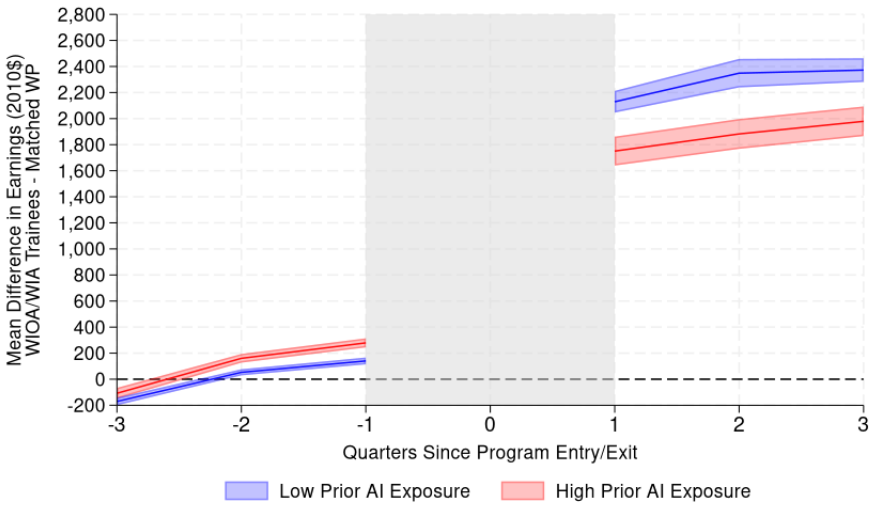
Figure 8: Common Keywords in AI Program Descriptions (Pre- vs. Post-2022)



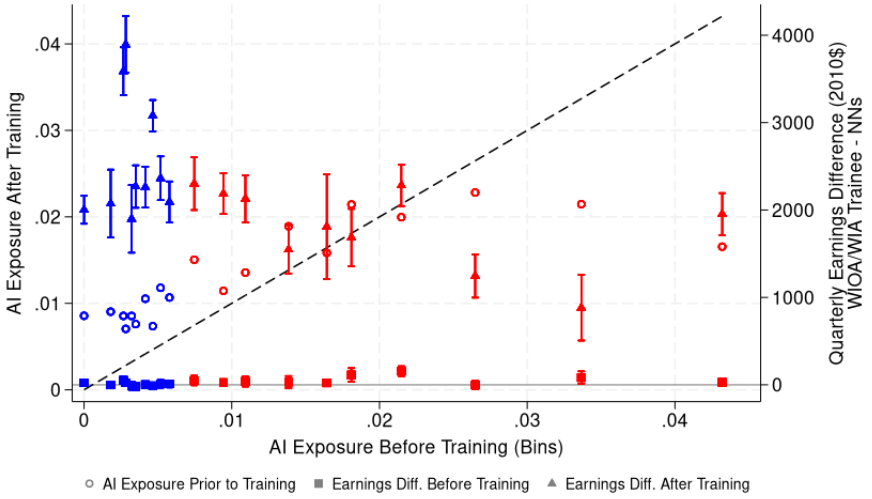
One final piece of evidence can be attained by estimating the earnings returns based on occupations which have been shown to use LLMs intensively in practice (rather than potentially, as captured by our baseline measure). Here, we use the Claude-based measure from Anthropic and reproduce our main estimates: (i) the overall returns to training, as well as the (ii) heterogeneous returns for different skill transitions. Figure 9 shows that the results based on exposure measures from Claude usage data are remarkably similar to what was found using the potential AI exposure measure. Taken together, these results suggest that using historical returns to cognitive training can be informative for more present-day, AI-related training.

Figure 9: NN-Matched Earnings Returns to Job Training: Claude-Based Prior Exposure

(A) Using 55th Percentile in Claude Use Cutoff for "High"



(B) Earnings Returns to Job Training vs. Skill Transitions: Claude-Based



Notes: AI exposure is based on the Anthropic Economic Index Claude-based occupation-level AI exposure scores (Handa et al., 2025). Each series in Panel A represents the earnings return relative to nearest-neighbor matched twins. Panel B overlays pooled estimates for the returns to training across the AI exposure spectrum (triangles), placebo estimates for pooled estimates prior to entry into training/job search services for the control group (squares), and average movements in AI skill before and after training (circles). Circles above the 45 degree line represent AI upskilling, while below the line represent AI downskilling. Red and blue correspond to our prior preferred high and low exposure distinction. The Claude use distribution skews left; one estimate to the right of the support is suppressed for exposition.

5 AI Retrainability (AIR) Index

We have established that the earnings returns to training are high across the AI spectrum, and that they are mediated by AI skill movements. In this section, we focus on the best-case

scenario: workers who jointly move up in AI skill space after job training, and receive higher earnings for doing so.

5.1 What share of workers and occupations are AI retrainable?

Using our nearest-neighbor matched estimates, we can calculate the number of training participants for whom both earnings returns and the AI content of their next occupation are higher than those of their matched pair counterparts. In Table A.6, we show the share of training spells for which both AI intensity and earnings are weakly greater, separately for the training group relative to matched nearest neighbors. We find that about 38% of workers are AI retrainable, with this estimate declining slightly to 37% when using the LLM measure developed by Handa et al. (2025). We also perform the same calculation by asking for any given 6-digit occupation, whether on average workers displaced from that occupation move to higher earnings and AI exposure scores, and find 45%-50% of occupations doing so. These AI retrainable occupations cover 45% of the CPS labor force, though this is likely a lower bound as some occupations in the CPS do not appear in WIOA/WIA.

5.2 AI Retrainability Index

To further unpack the mechanisms underlying the estimated large degree of adaptability to AI shocks through retraining, we develop an "AI Retrainability Index" (AIR) that ranks occupations based on the share of job training participants who receive positive returns from moving to more AI-intensive work, relative to their matched control group pairs. This allows us to decompose whether training success emanates from movements up the AI skill ladder, versus movements in earnings holding AI skills constant. Higher weight on the latter would be consistent with alternative stories unrelated to AI upskilling, such as overcoming occupational licensing barriers through job training. That is, here we seek to separately isolate the joint probability of "AI retrainability" from overall retrainability.

We use information from each nearest-neighbor matched pair p. Training participants (and their matched pairs) begin in origin occupation i prior to training and can flow to any destination occupation j after training.[18] The index is as follows

$$AIR_i = \sum_j \frac{1}{n_{ij}} \Delta AI_{ij}^{\sigma} * \mathcal{M}(\Delta \log(\overline{Y_{ij}})) \tag{1}$$

[18]In this descriptive exercise, we further strip out any participants who were dual enrolled in WIOA job training and Wagner-Peyser search assistance to hone in on the mechanisms driving the high returns to job training in particular.

where n_{ij} is the number of trainees flowing from occupation i to j, capturing probability of reemployment in occupations of varying AI intensities. $\Delta AI_{ij}^{\sigma} = (AI_j^{\sigma} - AI_i^{\sigma})_{p=1} - (AI_j^{\sigma} - AI_i^{\sigma})_{p=0}$ is the difference in standardized mean Eloundou et al. (2024) AI exposure scores associated with each trainee's occupational transition from i to j relative to the change in AI exposure score for their matched control group unit. $\Delta \log(\overline{Y_{ij}}) = \frac{1}{p_{ij}} \sum_p \Delta \log(Y_{p=1}^{i \to j}) - \Delta \log(Y_{p=0}^{i})$ is the mean returns to training for each treated unit in the nearest neighbor pair p relative to each control group unit $p{=}0$ that started in same occupation i but is free to move to any subsequent sector j. $\mathcal{M} \in [0, 1]$ is a MinMax transformation that allows us to circumvent the problem of multiplying negative wage returns by negative AI exposure changes by converting the most negative earnings returns into low positive values:

$$\mathcal{M}(\Delta \log(\overline{Y_{ij}})) = \frac{\Delta \log(\overline{Y_{ij}}) - \min(\Delta \log(\overline{Y_{ij}}))}{\max(\Delta \log(\overline{Y_{ij}})) - \min(\Delta \log(\overline{Y_{ij}}))} \tag{2}$$

The functional form in Equation 1 is deliberately chosen such that the index has the largest value when both the change in the AI content of work (i.e. AI exposure) and earnings returns to training are highly positive, and the index can only turn negative when the change in the AI content of work is negative. This allows us to hone in on whether the highest ranked occupations in AI retrainability are driven by much higher earnings gains or instead large leaps in AI skill after job training. By differencing out effects with respect to nearest neighbor pairs, this functional form also minimizes any mechanical bias that may arise from occupations with more or less scope to move in AI skill space due to starting at initially low or high SML values.[19]

To see this clearly, Figure 10 presents the overall 2-digit rankings of occupations by their AI retrainability. Panel (A) shows that only three occupations at the 2-digit level show positive improvements in AI skill (though this aggregation masks a larger number of 6-digit occupations that have positive AI skill returns). The top ranked 2-digit occupations include some expected occupations like *Legal*, however others that appear more manual in nature but whose trainees have more scope to retrain. The interpretation of the top ranked occupation is that workers displaced from legal professions (such as paralegals who comprise the largest disaggregated occupation in the *Legal* 2-digit category) are the most AI-retrainable workers in the index. Panel (B) decomposes the contribution of AI skill and earnings returns components to the index, where index values are expressed as a heatmap with top ranked occupations in darker reds. Here we can clearly see that higher heatmap (index) values are driven by moving from bottom to top along the AI skill dimension (as the

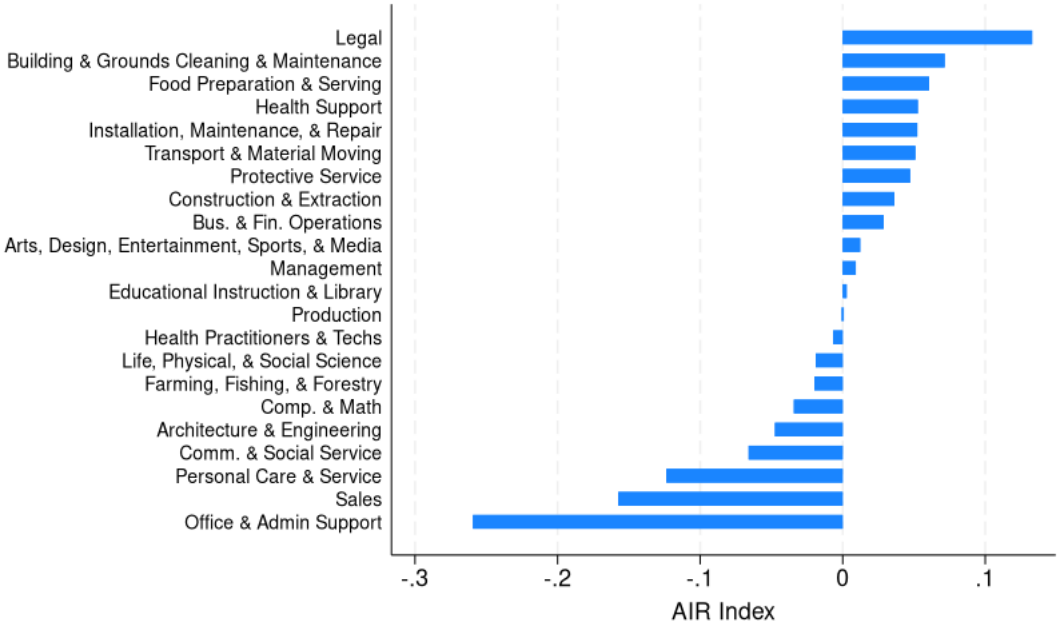
[19]In practice, the relative rankings of occupations remain relatively stable across a number of functional form choices for Equation 1, however this form clarifies that negative values correspond to negative movements in AI skill.

index intends), rather than from left to right in earnings returns. Panel (B) also reveals that more manual occupations have higher rankings, in part, beacuse of unconditonally higher earnings returns rather than higher earnings returns mediated by increases in AI skill (as is clearly the case with the *Legal* sector).

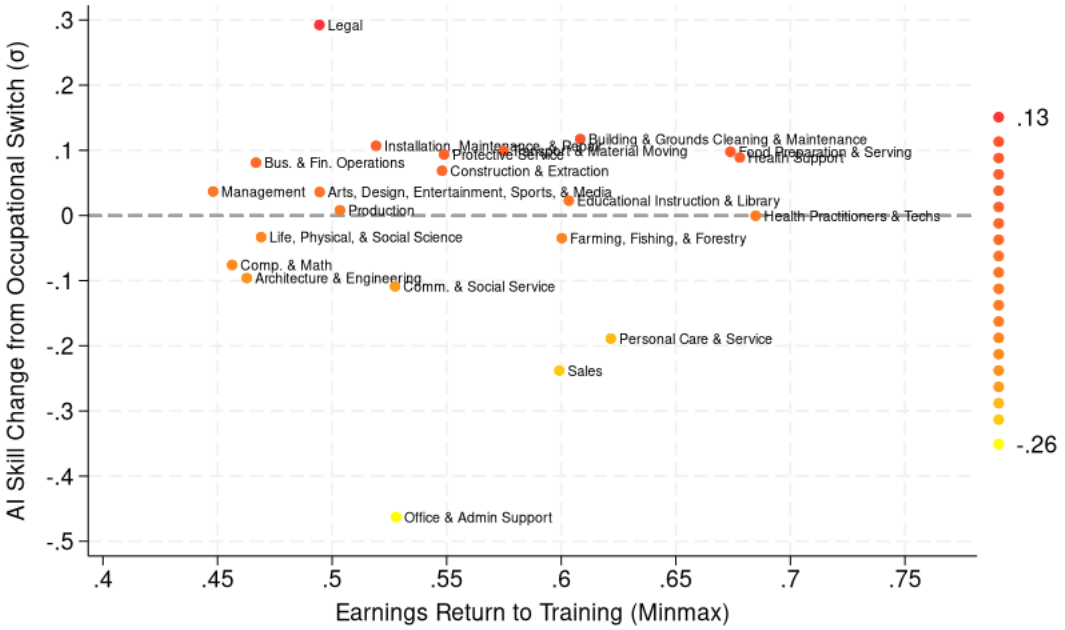
In a final exercise, Figure 11 shows decompositions for the modal 6-digit occupation within three illustrative occupations. The figure shows decompositions for the modal 6-digit occupation within three illustrative occupations: *Legal* (AIR Rank 1/23) whose dominant 6-digit occupation is *Paralegals*, *Health Support*(AIR Rank 4/23) whose dominant 6-digit occupation is *Nursing Assistants*, and *Office and Admin Support* (AIR Rank 23/23) whose dominant occupation is *Customer Service Representatives*. The figure shows that paralegals largely move up in AI skills after retraining, adapting to professions such as *Insurance Claims Clerks* which are presumably machine-substitutable. By contrast, *Nursing Assistants* move to higher earnings values while keeping AI skills constant, consistent with overcoming credentialing barriers to becoming *Registered Nurses* (also documented by Jacobson and Davis (2017)). Finally, *Customer Service Representatives* are not as adaptable, and instead find jobs in more physical professions that are insulated from AI, such as *Truck Drivers and Movers.* Overall, these decompositions show that occupations vary tremendously in their adaptability to AI, with some workers deepening their AI skills and others avoiding them.

Figure 10: AI Retrainability Index

(A) Index Rankings (2-Digit SOC)



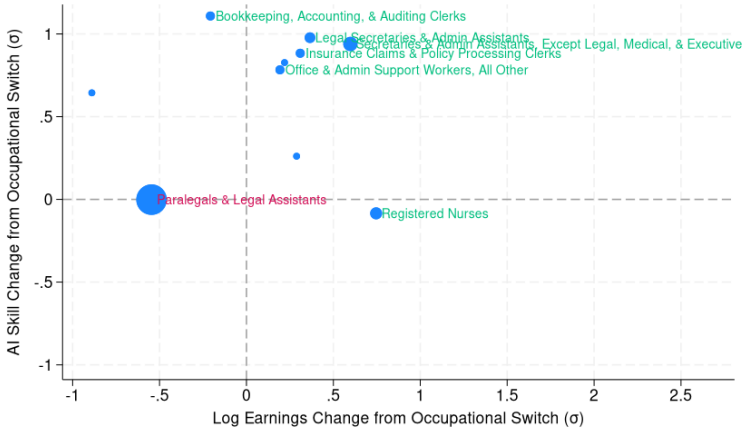
(B) Decomposition



Notes: Panel (A) shows 2-digit rankings of AI retrainability index (AIR) using Eloundou et al. (2024) AI exposure measure, where higher values reflect joint movements up in earnings AI skiill space. Panel (B) shows a decomposition with a heat map for higher index values, in which trainees leaving occupations above the horizontal dashed line are moving up in AI skill space.

Figure 11: AIR Index: Decomposition of Illustrative 6-Digit SOC Examples

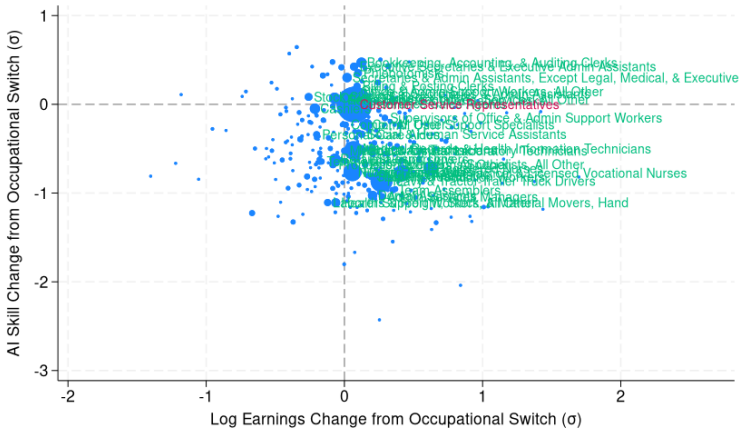
(A) Paralegals (Legal = 2-Digit AIR Rank 1/23)



(B) Nursing Assistants (Health Support = 2-Digit AIR Rank 4/23)



(C) Customer Service Representatives (Office & Admin Support = 2-Digit AIR Rank 23/23)



Notes: Figures show flows in AI intensity (y-axis) and earnings returns from training (x-axis) relative to nearest neighbor matched pairs, for three origin occupations (prior to program participation) along with destination occupations. Red text indicates flows within the same 6-digit occupation.

6 Discussion

One of the most pressing policy questions today is the effects that AI technologies will have in the labor market, specifically their effects on workers. One way that workers can respond to these new technologies, thus far understudied, is to retrain themselves - to either reskill to use the technology, or reskill to avoid it. In the context of a virtually nonexistent AI re-skilling literature, our main contribution is to provide a national and large-scale lens into AI reskilling efforts that goes beyond individual firm training studies.

How retrainable are AI-exposed workers? We find that workers in AI-exposed occupations are surprisingly resilient in adjusting to AI pressures through job training. Comparing high AI exposure trainees to a matched sample of similar workers who only received job search assistance, we find that AI-exposed workers have high earnings returns from training ($\$1,400$ per quarter), although they are lower than trainees from low exposure occupations ($\$2,600$ per quarter). Trainees from low-exposure occupations experience gains by upskilling in AI space, while those from high-exposure occupations experience gains by downskilling in AI space, suggesting that a worker's optimal AI skill investment may lie in an intermediate range.

Using a newly constructed AI retrainability index (AIR), we estimate that about 45-50% of occupations are "AI-retrainable"—workers leaving these occupations are compensated with higher earnings from moving up the AI skill ladder – a large share given our relatively low-income sample of displaced workers. Our evidence thus indicates that with appropriate investment in training infrastructure, workers can successfully adapt to AI technologies, lending support to more optimistic scenarios in which policy interventions enable adjustment to AI instead of displacement by it.

References

Abadie, Alberto and Guido Imbens, "Simple and bias-corrected matching estimators for average treatment effects," 2002.

_ , David Drukker, Jane Leber Herr, and Guido W Imbens, "Implementing matching estimators for average treatment effects in Stata," *The stata journal*, 2004, *4* (3), 290–311.

Abel, Jaison R, Richard Deitz, Natalia Emanuel, and Benjamin Hyman, "AI and the Labor Market: Will Firms Hire, Fire, or Retrain?," *Liberty Street Economics*, 2024.

Acemoglu, Daron and David Autor, "Skills, tasks and technologies: Implications for employment and earnings," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1043–1171.

Andersson, Fredrik, Harry J Holzer, Julia I Lane, David Rosenblum, and Jeffrey Smith, "Does Federally Funded Job Training Work?: Nonexperimental Estimates of WIA Training Impacts Using Longitudinal Data on Workers and Firms," *Journal of Human Resources*, 2024, *59* (4), 1244–1283.

Ashenfelter, Orley, "Estimating the Effect of Training Programs on Earnings," in Orley Ashenfelter and James Blum, eds., *Evaluating the Labor Market Effects of Social Programs*, Princeton University Press, 1978, pp. 95–124.

Autor, David, "Applying AI to rebuild middle class jobs," Technical Report, National Bureau of Economic Research 2024.

Bollinger, Christopher R. and Kenneth R. Troske, "Evaluation of a New Job Training Program: Code Louisville," *Journal of Policy Analysis and Management*, 2025, —.

Bonney, Kathryn, Cory Breaux, Cathy Buffington, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff, and Keith Savage, "Tracking firm use of AI in real time: A snapshot from the Business Trends and Outlook Survey," Technical Report, National Bureau of Economic Research 2024.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," Working Paper, Stanford Digital Economy Lab 2025. Accessed 1 Feb 2026.

_ , Danielle Li, and Lindsey Raymond, "Generative AI at work," *The Quarterly Journal of Economics*, 2025, p. qjae044.

_ , Tom Mitchell, and Daniel Rock, "What can machines learn and what does it mean for occupations and the economy?," in "AEA papers and proceedings," Vol. 108 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2018, pp. 43–47.

Card, David, Jochen Kluve, and Andrea Weber, "What works? A meta analysis of recent active labor market program evaluations," *Journal of the European Economic Association*, 2018, *16* (3), 894–931.

Delfino, Alexia, Andrea Garnero, Sergio Inferrera, Marco Leonardi, and Raffaella Sadun, "Unwilling to Reskill? Experimental Evidence from Real-World Jobseekers," Technical Report w34633, National Bureau of Economic Research 2026.

Deming, David J, Christopher Ong, and Lawrence H Summers, "Technological disruption in the labor market," Technical Report, National Bureau of Economic Research 2025.

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

Hampole, Menaka, Dimitris Papanikolaou, Lawrence Schmidt, and Bryan Seegmiller, "Artificial Intelligence and the Labor Market," *Available at SSRN*, 2025.

Handa, Kunal, Alex Tamkin, Miles McCain, Saffron Huang, Esin Durmus, Sarah Heck,

Jared Mueller, Jerry Hong, Stuart Ritchie, Tim Belonax et al., "Which economic tasks are performed with ai? evidence from millions of claude conversations," *arXiv preprint arXiv:2503.04761*, 2025.

Hartley, Jonathan, Filip Jolevski, Vitor Melo, and Brendan Moore, "The Labor Market Effects of Generative Artificial Intelligence," *Available at SSRN*, 2024.

Heckman, James J., Hidehiko Ichimura, Jeffrey A. Smith, and Petra E. Todd, "Characterizing Selection Bias Using Experimental Data," *Econometrica*, 1998, *66* (5), 1017–1098.

Humlum, Anders and Emilie Vestergaard, "The adoption of ChatGPT," Technical Report, IZA Discussion Papers 2024.

_ , Jakob Munch, and Mette Rasmussen, "What works for the unemployed? Evidence from quasi-random caseworker assignments," 2023.

Hyman, Ben, Ben Lahey, Karen Ni, and Laura Pilossoph, "Job Training Mismatch," 2025.

Jacobson, Louis and Jonathan Davis, "The relative returns to workforce investment act–supported training in florida by field, gender, and education and ways to improve trainees' choices," *Journal of Labor Economics*, 2017, *35* (S1), S337–S375.

Lechner, Michael, Ruth Miquel, and Conny Wunsch, "Long-Run Effects of Public Sector Sponsored Training in West Germany," *Journal of the European Economic Association*, 2011, *9* (4), 742–784.

Pizzinelli, Carlo, Augustus J Panton, Ms Marina Mendes Tavares, Mauro Cazzaniga, and Longji Li, *Labor market exposure to AI: Cross-country differences and distributional implications*, International Monetary Fund, 2023.

Rothstein, Jesse, Robert Santillano, Till Von Wachter, Wahid Khan, and Mary Yang, "CAAL-Skills: Study of Workforce Training Programs in California," 2022.

Şahin, Ayşegül, Joseph Song, Giorgio Topa, and Giovanni L Violante, "Mismatch unemployment," *American Economic Review*, 2014, *104* (11), 3529–3564.

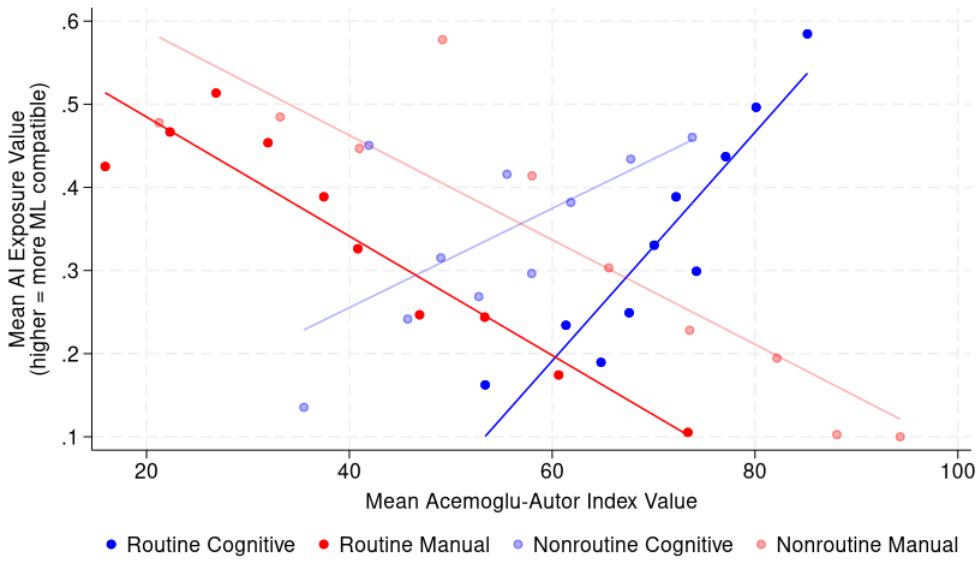
Tomlinson, Kiran, Sonia Jaffe, Will Wang, Scott Counts, and Siddharth Suri, "Working with AI: Measuring the Occupational Implications of Generative AI," *arXiv preprint arXiv:2507.07935*, 2025.

Appendices

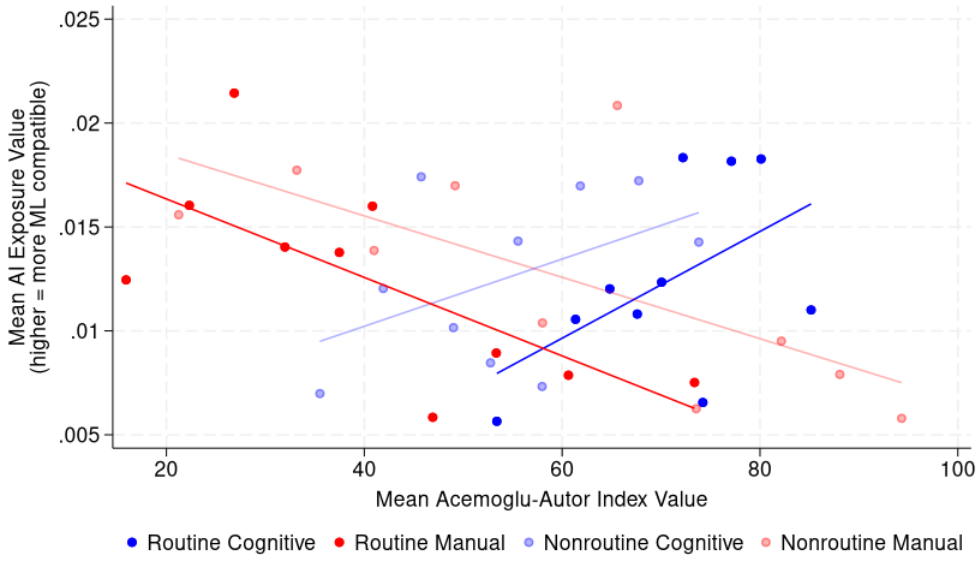
Appendix A. Additional Results

Figure A.1: AI Exposure Correlations with Acemoglu and Autor (2011) Task Measures

(A) Eloundou AI Exposure v. AA Index

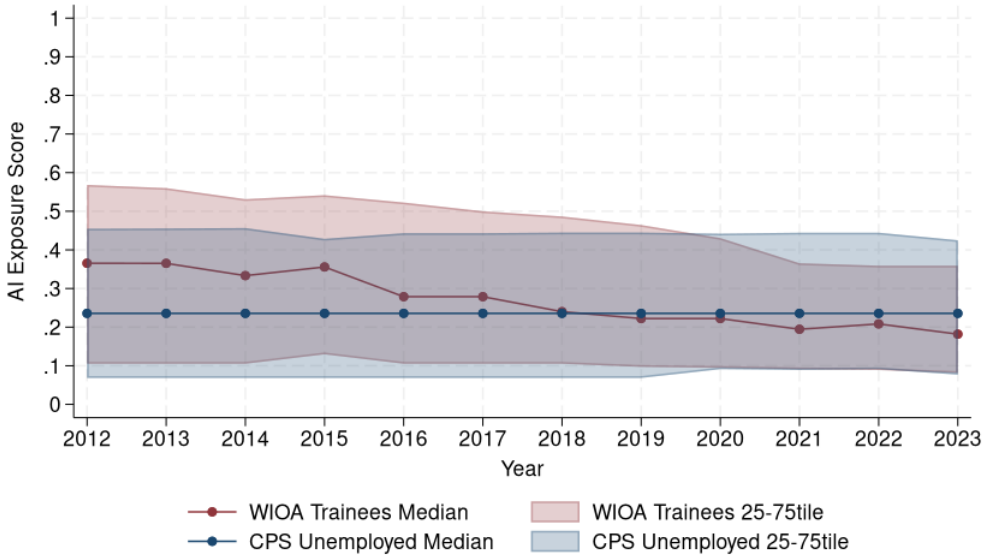


(B) Claude AI Exposure vs. AA Index



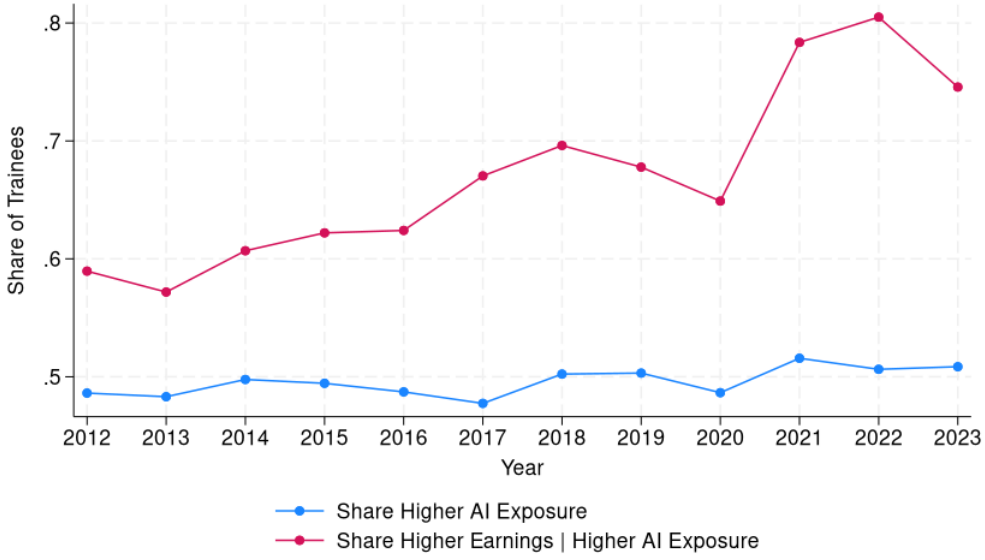
Notes: The x-axis in both panels measures the mean Acemoglu and Autor (2011) (AA) task share for each 6-digit occupation. The y-axis in panel (A) is our main AI exposure measure—-Eloundou et al. (2024) at the 6-digit occupation level. The y-axis for panel (B) represents the secondary AI exposure measure based on Claude task usage from Anthropic's Econoimc Index (Handa et al., 2025). Correlations between AA task shares and AI exposure measures are expressed as binscatter plots.

Figure A.2: WIOA and CPS-Unemployed AI Exposure Distributions Over Time



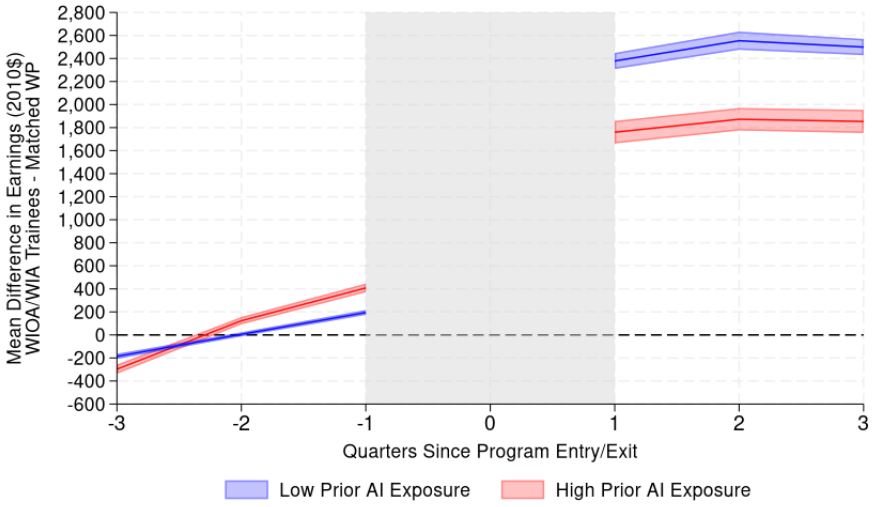
Notes: Figure plots the inter-quartile range of occupation AI exposure scores (Eloundou et al. (2024)) for WIOA/WIA trainees' occupations prior to training, and CPS unemployed workers' occupations prior to their current unemployment spell.

Figure A.3: Share of Trainees With Positive AI Skill and Earnings Changes over Time



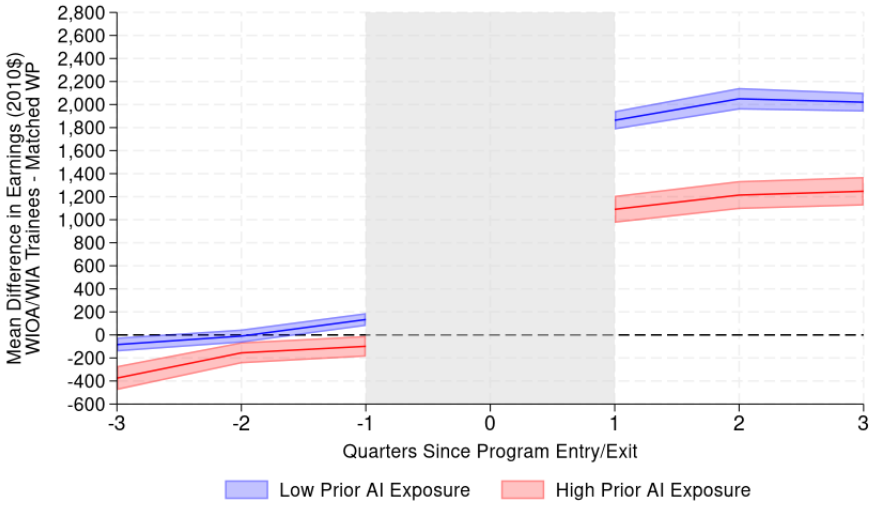
Notes: Figure plots share of WIOA/WIA training participants moving to higher AI exposure occupations (Eloundou et al. (2024)) after training by calendar year of participation, and jointly receiving higher earnings from moving to higher AI exposure values.

Figure A.4: Nearest-Neighbor Matched Returns to WIOA/WIA: Pooled Across Years: Robustness to Including Zero Earnings



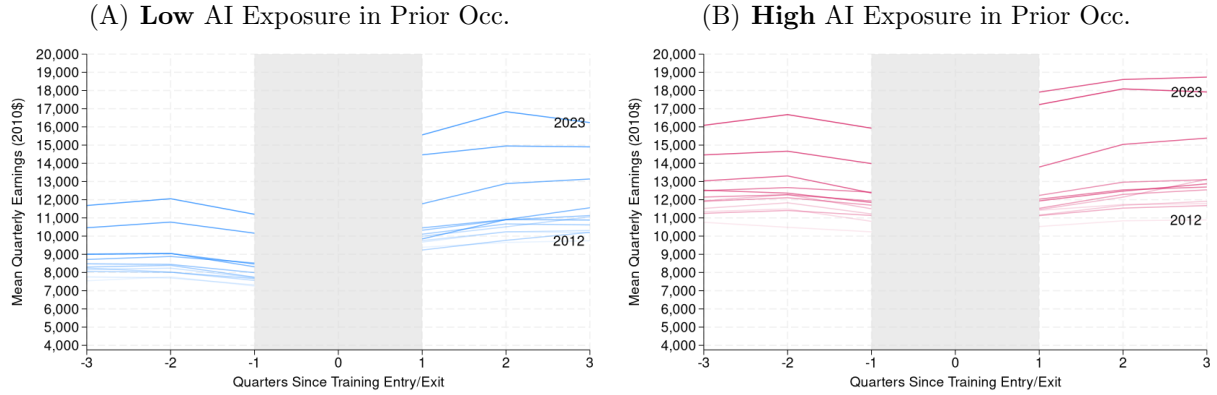
Notes: Figure is equivalent to main results Figure 3, but includes observations in which earnings are recorded as zero.

Figure A.5: Nearest-Neighbor Earnings Returns to Training Earnings: Robustness to Dropping Prior Earnings Covariates from Matching



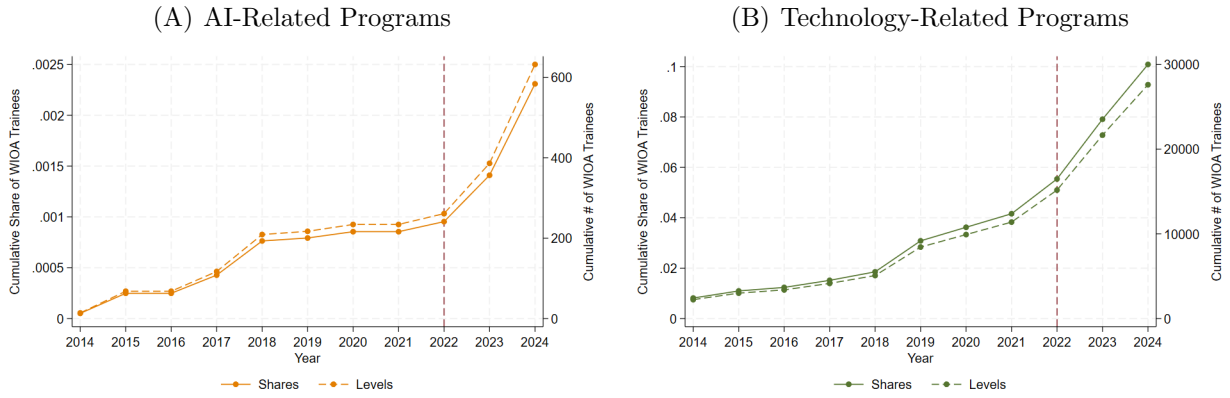
Notes: Figure is equivalent to main results Figure 3, but does not include any prior earnings variables in fuzzy or exact matching covariate lists.

Figure A.6: Earnings by AI Exposure Prior to Job Training, NN-Matched Control Group



Notes: Plots are computed for Wagner-Peyser matched nearest neighbors. Participants are classified as "high" AI exposure if their 6-digit SOC exposure score is above the 55th percentile of the Eloundou et al. (2024) 6-digit AI exposure score distribution, and "low" if below the 55th percentile. Heavier lines correspond to more recent years. The shaded gray region emphasizes that quarterly earnings are observed in the quarters prior to entry, and after exit, but not during program receipt. All earnings are in real 2010 dollars.

Figure A.7: Cumulative Growth in U.S. AI- and Technology-Related Job Training Participants



Notes: Panel A shows AI-related programs (e.g. keywords: *Artificial Intelligence, Natural Language Processing, Machine Learning*), and Panel B shows non-AI general Technology-related programs (e.g. keywords: *Programming, Web Development, Python*) drawn from WIOA ETPL lists over time, and weighted by the number of WIOA job training participants enrolled. Shares are calculated as the cumulative number of AI or Technology programs (trainee-weighted) introduced over the cumulative number of all WIOA ETPL programs (trainee-weighted) within a given year. See Appendix C for details.

Table A.1: Summary Statistics by AI Exposure Prior to Training Participation

	High Prior AI Exposure		Low Prior AI Exposure		High AI - Low AI	
	Mean/SD	# Training Spells	Mean/SD	# Training Spells	Δ/SE	% Diff
	(1)	(2)	(3)	(4)	(5)	(6)
A. Demographic						
Female	0.61	35,954	0.45	59,949	0.17	27.2
	[0.49]		[0.50]		(0.0033)	
Age	41.9	35,954	36.6	59,949	5.38	12.8
	[11.1]		[10.4]		(0.073)	
Asian	0.039	35,954	0.021	59,949	0.018	45.9
	[0.19]		[0.14]		(0.0011)	
Black	0.23	35,954	0.32	59,949	-0.086	-37.3
	[0.42]		[0.47]		(0.0030)	
White	0.62	35,954	0.56	59,949	0.059	9.53
	[0.49]		[0.50]		(0.0033)	
Hispanic/Latino	0.15	35,954	0.15	59,949	0.0027	1.80
	[0.36]		[0.36]		(0.0024)	
Disability Status	0.027	35,954	0.022	59,949	0.0046	17.4
	[0.16]		[0.15]		(0.00098)	
No HS Diploma/GED	0.022	35,954	0.056	59,949	-0.033	-149.4
	[0.15]		[0.23]		(0.0012)	
HS Diploma/GED	0.37	35,954	0.56	59,949	-0.20	-53.3
	[0.48]		[0.50]		(0.0033)	
Some College	0.20	35,954	0.19	59,949	0.0069	3.52
	[0.40]		[0.39]		(0.0027)	
College Degree Plus	0.41	35,954	0.19	59,949	0.22	54.4
	[0.49]		[0.39]		(0.0030)	
B. Social Benefits						
Disadvantaged Adult Status	0.38	35,954	0.64	59,949	-0.26	-67.8
	[0.49]		[0.48]		(0.0032)	
Dislocated Worker Status	0.62	35,954	0.35	59,949	0.27	44.1
	[0.49]		[0.48]		(0.0032)	
Low-income Status	0.37	35,954	0.50	59,949	-0.12	-32.6
	[0.48]		[0.50]		(0.0033)	
TANF Recipient	0.0058	35,954	0.0077	59,949	-0.0019	-32.8
	[0.076]		[0.087]		(0.00048)	
SNAP Recipient	0.17	25,546	0.25	48,736	-0.079	-45.5
	[0.38]		[0.44]		(0.0031)	
Other Public Assistance Recipient	0.21	27,320	0.29	52,728	-0.077	-36.5
	[0.41]		[0.45]		(0.0032)	
C. Pre-Separation Employment						
Employed at Participation	0.22	35,954	0.42	59,949	-0.20	-90.2
	[0.42]		[0.49]		(0.0030)	
Real Earnings 1Q Before	11899.7	35,954	8152.5	59,949	3747.1	31.5
	[9742.0]		[6263.4]		(57.1)	
Real Earnings 2Q Before	12481.3	35,954	8819.6	59,949	3661.8	29.3
	[8820.6]		[6004.0]		(52.2)	
Real Earnings 3Q Before	12341.0	35,954	8681.3	59,949	3659.7	29.7
	[8641.5]		[5939.1]		(51.4)	
Quarters in WIOA/WIA	5.14	35,954	4.53	59,949	0.61	12.0
	[3.41]		[3.24]		(0.022)	
Observations	35,954		59,949			

Notes: Table presents means, standard deviations, and sample counts for main sample of job training participants by High and Low AI exposure prior to job training participation (see text for details). The pooled sample of trainees contains 95,903 participation spells. T-test standard errors in column (5) are adjusted for unequal variances (heteroskedasticity) across groups. Participants are weighted by the inverse distance to their nearest-neighbor from Wagner-Peyser used for results estimation.

Table A.2: Attrition Balance for WIOA/WIA Trainees with Missing Occupation Codes

	Nonmissing Occ Codes		Missing Occ Codes		Nonmiss - Miss	
	Mean/SD	# Training Spells	Mean/SD	# Training Spells	Δ/SE	% Diff
	(1)	(2)	(3)	(4)	(5)	(6)
A. Demographic						
Female	0.51	95,903	0.51	484,754	-0.000062	-0.012
	[0.50]		[0.50]		(0.0018)	
Age	38.6	95,903	38.0	403,882	0.62	1.60
	[11.0]		[11.0]		(0.039)	
Asian	0.041	76,452	0.035	416,124	0.0057	14.0
	[0.20]		[0.18]		(0.00073)	
Black	0.34	79,723	0.31	420,272	0.021	6.39
	[0.47]		[0.46]		(0.0018)	
White	0.66	83,093	0.64	434,872	0.028	4.16
	[0.47]		[0.48]		(0.0018)	
Hispanic/Latino	0.17	91,447	0.16	439,054	0.0052	3.17
	[0.37]		[0.37]		(0.0013)	
Disability Status	0.028	93,590	0.028	457,672	0.00051	1.79
	[0.17]		[0.16]		(0.00059)	
No HS Diploma/GED	0.046	95,903	0.058	484,754	-0.012	-25.8
	[0.21]		[0.23]		(0.00081)	
HS Diploma/GED	0.49	95,903	0.47	484,754	0.019	3.92
	[0.50]		[0.50]		(0.0018)	
Some College	0.19	95,903	0.19	484,754	0.0021	1.09
	[0.39]		[0.39]		(0.0014)	
College Degree Plus	0.27	95,903	0.28	484,754	-0.0060	-2.17
	[0.45]		[0.45]		(0.0016)	
B. Social Benefits						
Disadvantaged Adult Status	0.55	95,903	0.65	484,754	-0.10	-19.1
	[0.50]		[0.48]		(0.0017)	
Dislocated Worker Status	0.45	95,903	0.37	484,754	0.077	17.2
	[0.50]		[0.48]		(0.0017)	
Low-income Status	0.45	95,903	0.43	453,209	0.020	4.45
	[0.50]		[0.50]		(0.0018)	
TANF Recipient	0.011	84,489	0.013	406,762	-0.0027	-24.6
	[0.10]		[0.12]		(0.00043)	
SNAP Recipient	0.23	74,282	0.22	252,760	0.011	4.78
	[0.42]		[0.41]		(0.0017)	
Other Public Assistance Recipient	0.26	80,048	0.22	385,721	0.041	15.5
	[0.44]		[0.42]		(0.0016)	
C. Pre-Separation Employment						
Employed at Participation	0.34	95,903	0.41	484,754	-0.062	-17.9
	[0.48]		[0.49]		(0.0017)	
Real Earnings 1Q Before	9611.7	95,903	9754.8	484,754	-143.1	-1.49
	[8062.5]		[13326.9]		(44.6)	
Real Earnings 2Q Before	10267.5	95,903	10224.4	484,754	43.1	0.42
	[7508.3]		[12338.6]		(41.3)	
Real Earnings 3Q Before	10124.8	95,903	9993.1	484,754	131.7	1.30
	[7396.2]		[11141.1]		(37.5)	
Quarters in WIOA/WIA	4.76	95,903	4.38	484,754	0.38	8.07
	[3.32]		[3.15]		(0.011)	
Observations	95,903		484,754			

Notes: Columns (1) and (2) show summary statistics for the main analysis sample of job trainees (N=95,903 training spells). Columns (3) and (4) present summary statistics for training spells with partially missing occupation data either prior to training entry or after training exit. T-test standard errors in column (5) are adjusted for unequal variances (heteroskedasticity) across groups.

Table A.3: Balance Table: Trainees vs. Matched WP Neighbors (Low Initial AI)

	Trainees		Matched Wagner Peyser		Trainees - Wagner Peyser	
	Mean/SD	# Spells	Mean/SD	# Spells	Δ/SE	% Diff
	(1)	(2)	(3)	(4)	(5)	(6)
A. Demographic						
Female	0.45	59,949	0.47	59,949	-0.025	-5.48
	[0.50]		[0.50]		(0.0029)	
Age	36.6	59,949	38.2	59,949	-1.62	-4.43
	[10.4]		[10.8]		(0.061)	
Asian	0.021	59,949	0.0076	59,949	0.013	63.8
	[0.14]		[0.087]		(0.00063)	
Black	0.32	59,949	0.47	59,949	-0.15	-47.8
	[0.47]		[0.50]		(0.0028)	
White	0.56	59,949	0.48	59,949	0.082	14.6
	[0.50]		[0.50]		(0.0029)	
Hispanic/Latino	0.15	59,949	0.074	59,949	0.076	50.6
	[0.36]		[0.26]		(0.0018)	
Disability Status	0.022	59,949	0.0094	59,949	0.013	57.3
	[0.15]		[0.096]		(0.00067)	
No HS Diploma/GED	0.056	59,949	0.038	59,949	0.018	31.9
	[0.23]		[0.19]		(0.0012)	
HS Diploma/GED	0.56	59,949	0.64	59,949	-0.071	-12.5
	[0.50]		[0.48]		(0.0028)	
Some College	0.19	59,949	0.19	59,949	0.0035	1.83
	[0.39]		[0.39]		(0.0023)	
College Degree Plus	0.19	59,949	0.13	59,949	0.052	27.9
	[0.39]		[0.34]		(0.0021)	
B. Social Benefits						
Disadvantaged Adult Status	0.64	59,949	0.39	59,949	0.25	39.4
	[0.48]		[0.49]		(0.0028)	
Dislocated Worker Status	0.35	59,949	0.27	59,949	0.081	23.4
	[0.48]		[0.44]		(0.0027)	
Low-income Status	0.50	59,949	0.41	59,949	0.090	18.0
	[0.50]		[0.49]		(0.0029)	
TANF Recipient	0.0077	59,949	0.0028	59,949	0.0049	63.1
	[0.087]		[0.053]		(0.00038)	
SNAP Recipient	0.25	48,736	0.14	49,232	0.11	44.6
	[0.44]		[0.35]		(0.0025)	
Other Public Assistance Recipient	0.29	52,728	0.28	33,374	0.0066	2.30
	[0.45]		[0.45]		(0.0032)	
C. Pre-Separation Employment						
Employed at Participation	0.42	59,949	0.31	59,949	0.11	25.9
	[0.49]		[0.46]		(0.0028)	
Real Earnings 1Q Before	8152.5	59,949	8343.8	59,949	-191.2	-2.35
	[6263.4]		[5965.5]		(35.0)	
Real Earnings 2Q Before	8819.6	59,949	8790.7	59,949	28.9	0.33
	[6004.0]		[5801.2]		(33.8)	
Real Earnings 3Q Before	8681.3	59,949	8513.5	59,949	167.8	1.93
	[5939.1]		[5675.0]		(33.3)	
Quarters in WIOA/WP	4.53	59,949	2.56	59,949	1.97	43.4
	[3.24]		[2.40]		(0.017)	

Notes: Columns (1) through (4) show summary statistics for low initial AI exposure trainees (using Eloundou et al. (2024)) and nearest neighbor maches, while columns (5) and (6) estimate t-tests and standard errors adjusted for unequal variances (heteroskedasticity) across groups. Participants are weighted by the inverse distance to their nearest-neighbor from Wagner-Peyser used for results estimation.

Table A.4: Balance Table: Trainees vs. Matched WP Neighbors (High Initial AI)

	Trainees		Matched Wagner Peyser		Trainees - Wagner Peyser	
	Mean/SD	# Spells	Mean/SD	# Spells	Δ/SE	% Diff
	(1)	(2)	(3)	(4)	(5)	(6)
A. Demographic						
Female	0.61	35,954	0.67	35,954	-0.051	-8.24
	[0.49]		[0.47]		(0.0036)	
Age	41.9	35,954	43.7	35,954	-1.80	-4.30
	[11.1]		[11.0]		(0.083)	
Asian	0.039	35,954	0.022	35,954	0.017	44.0
	[0.19]		[0.15]		(0.0012)	
Black	0.23	35,954	0.24	35,954	-0.0073	-3.14
	[0.42]		[0.43]		(0.0032)	
White	0.62	35,954	0.67	35,954	-0.054	-8.66
	[0.49]		[0.47]		(0.0036)	
Hispanic/Latino	0.15	35,954	0.100	35,954	0.053	34.9
	[0.36]		[0.30]		(0.0025)	
Disability Status	0.027	35,954	0.016	35,954	0.010	38.0
	[0.16]		[0.13]		(0.0010)	
No HS Diploma/GED	0.022	35,954	0.013	35,954	0.0097	43.4
	[0.15]		[0.11]		(0.00093)	
HS Diploma/GED	0.37	35,954	0.38	35,954	-0.012	-3.15
	[0.48]		[0.49]		(0.0036)	
Some College	0.20	35,954	0.19	35,954	0.011	5.55
	[0.40]		[0.39]		(0.0030)	
College Degree Plus	0.41	35,954	0.42	35,954	-0.0089	-2.17
	[0.49]		[0.49]		(0.0037)	
B. Social Benefits						
Disadvantaged Adult Status	0.38	35,954	0.28	35,954	0.11	27.6
	[0.49]		[0.45]		(0.0035)	
Dislocated Worker Status	0.62	35,954	0.60	35,954	0.024	3.93
	[0.49]		[0.49]		(0.0037)	
Low-income Status	0.37	35,954	0.31	35,954	0.063	16.9
	[0.48]		[0.46]		(0.0035)	
TANF Recipient	0.0058	35,954	0.0033	35,954	0.0025	42.6
	[0.076]		[0.058]		(0.00047)	
SNAP Recipient	0.17	25,546	0.13	25,765	0.045	26.0
	[0.38]		[0.34]		(0.0032)	
Other Public Assistance Recipient	0.21	27,320	0.19	22,581	0.019	9.22
	[0.41]		[0.39]		(0.0036)	
C. Pre-Separation Employment						
Employed at Participation	0.22	35,954	0.15	35,954	0.073	32.8
	[0.42]		[0.36]		(0.0029)	
Real Earnings 1Q Before	11899.7	35,954	12019.6	35,954	-120.0	-1.01
	[9742.0]		[9092.9]		(70.0)	
Real Earnings 2Q Before	12481.3	35,954	12322.7	35,954	158.7	1.27
	[8820.6]		[8338.1]		(63.7)	
Real Earnings 3Q Before	12341.0	35,954	12147.9	35,954	193.1	1.56
	[8641.5]		[8236.6]		(62.8)	
Quarters in WIOA/WP	5.14	35,954	3.15	35,954	2.00	38.8
	[3.41]		[2.48]		(0.022)	

Notes: Columns (1) through (4) show summary statistics for high initial AI exposure trainees (using Eloundou et al. (2024)) and nearest neighbor maches, while columns (5) and (6) estimate t-tests and standard errors adjusted for unequal variances (heteroskedasticity) across groups. Participants are weighted by the inverse distance to their nearest-neighbor from Wagner-Peyser used for results estimation.

Table A.5: Regression Estimates for Quarterly Earnings Returns to WIOA/WIA Training

	Prior AI Exposure (Eloundou)		**Prior AI Exposure (Claude)**	
	High	Low	High	Low
A. Pooled Post:	2,615	1,426	2,514	1,755
	(37.81)	(57.69)	(41.34)	(53.97)
B. Pooled Pre (Placebo):	77	2	74	10
	(10.34)	(3.52)	(9.23)	(4.23)
N	35,954	59,949	40,599	50,485
Mean Trainee Baseline Wages:	12,299	8,623	11,403	9050

Notes: Each cell reports a separate mean quarterly earnings regression estimate. All estimates are weighted by the inverse of the Mahalanobis distance for each nearest-neighbor matched pair. Mean baseline earnings reflect average quarterly earnings pooled across three quarters prior to WIOA/WIA participation. Prior AI exposure is measured prior to program enrollment, while target AI exposure is the occupation targeted or desired after training—see text for further details. Heteroskedasticity-robust standard errors are reported in parentheses.

Table A.6: Headline Results for Share of Participants and Occupations "AI-Retrainable"

		Trainee - Matched	
AI Retrainability Statistic		Full Sample (1)	High Initial AI Exposure (2)
Sh. Spells	$\Delta Eloundou \geq 0$ & $\Delta Y \geq 0$	38.0%	29.7%
Sh. Spells	$\Delta Claude \geq 0$ & $\Delta Y \geq 0$	36.5%	30.4%
Sh. Occups	$\Delta \overline{Eloundou} \geq 0$ & $\Delta \overline{Y} \geq 0$	45.3%	24.8%
Sh. Occups	$\Delta \overline{Claude} \geq 0$ & $\Delta \overline{Y} \geq 0$	49.7%	36.7%

Notes: Table shows counts of shares of workers and occupations for whom both the earnings returns from job training and the AI content of their subsequent occupation are higher after participation. Columns (1) and (2) show these as raw levels, whereas columns (3) and (4) show our preferred measures that calculate the number of workers for whom both outcomes are higher relative to the outcomes for matched nearest neighbor pairs. A 6-digit occupation is labeled AI-retrainable if the mean worker in that occupations AI-retrainable. See text for further details.

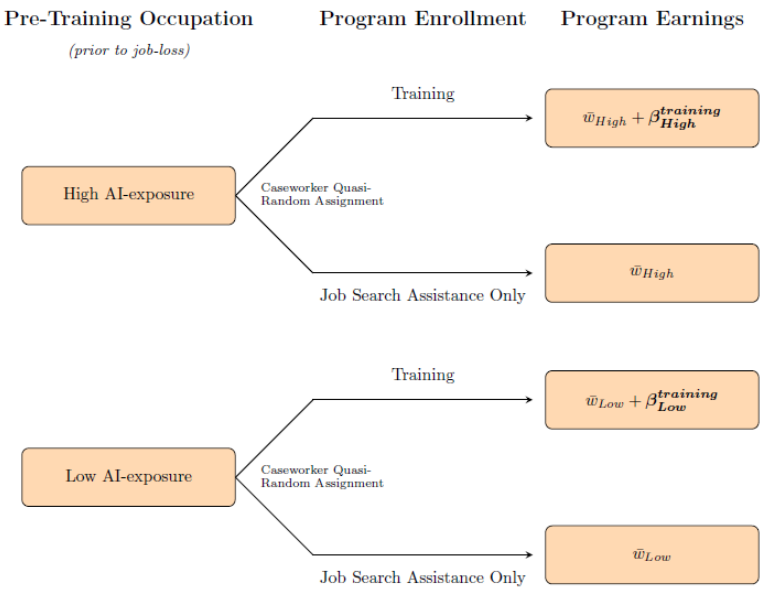
Table A.7: Top Occupations Prior to Training by AI Exposure (Claude-Use Based)

Rank	6-Digit SOC	Occupation	No. Trainees
AI Exposure Quintile 5 (High Exposure)			
1	412011	Cashiers	1,943
2	439061	Office clerks, general	1,377
3	412031	Retail salesperson	1,360
AI Exposure Quintile 4			
1	434051	Customer service representatives	3,929
2	119199	Managers, all other	1,206
3	533031	Driver/sales workers	853
AI Exposure Quintile 3			
1	533032	Heavy and tractor-trailer truck drivers	1,914
2	435081	Stockers and order filers	1,283
3	353031	Waiters and waitresses	982
AI Exposure Quintile 2			
1	311014	Nursing Assistants	4,981
2	537062	Laborers and freight, stock, and material movers, hand	1,923
3	519198	Helpers - production workers	1,549
AI Exposure Quintile 1 (Low Exposure)			
1	519199	Production workers, all other	2,564
2	533033	Light truck drivers	1,294
3	513092	Food batchmakers	1,252

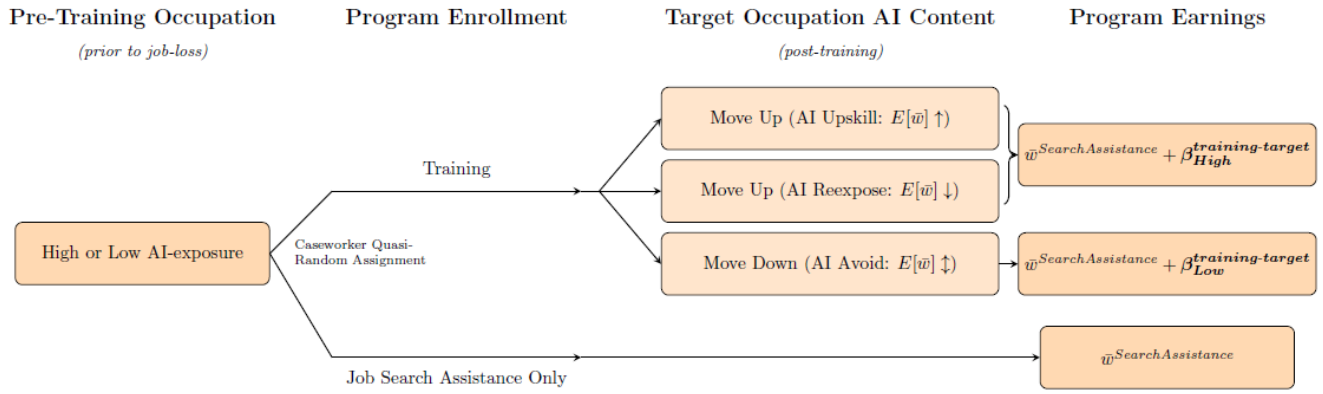
Notes: This table shows the top three most prominent occupations in different AI exposure quintiles prior to entering WIOA/WIA training participation. Counts of number of training participants are shown in the final column. Table uses AI exposure measures from Anthropic's Economic Index (Handa et al., 2025). See text for details.

Appendix B. Diagram Representations of Research Design

Main Design. Heterogeneous Design by Prior Exposure



Mediation Design. Heterogeneous Design by Skill Transition (Choice of Training)



Notes: Within the main design framework, returns to job training are estimated as a simple comparison of the WIOA program returns between trainees and their nearest neighbors who only receive job search assistance, producing $\beta^{training}_{High}$ and $\beta^{training}_{Low}$ for the two pre-training subgroups. Within the mediation design framework, returns to job training estimation requires a more demanding identification assumption that the nearest neighbor matched to trainees within each of the three target sub-branches serves as a good counterfactual for occupational transitions. That is, we need to assume WIOA trainees would have chosen similar occupations in the post-period as their matched Wagner Peyser counterparts in the absence of training.

Appendix C. WIOA ETPL Training Content Classification

To better understand how workforce development training programs evolve in their AI and technology content over time, we utilize the U.S. Department of Labor's comprehensive database of WIOA Eligible Training Provider Lists (ETPL), which contain detailed curricular and performance information on approved WIOA training programs (as reported by individual states on Form ETA-9171). The database covers training programs approved to be included in the WIOA ETPLs between January 2014 and October 2024, totaling 77,085 training programs across that time frame.

We utilize textual program descriptions reported by training providers to characterize each training program within this data as having (1) AI-specific content or (2) general technology content. To do this, we first fed a random sample of ∼200 training program descriptions from the WIOA ETPLs to an NLP tool which helped us identify common keywords and phrases related to AI or general technology training content. We then went through the comprehensive list of programs and tagged the relevant keywords within their program descriptions, creating an AI-specific tag and general Technology tag[20]. The keywords used for this tagging exercise are reported in Appendix Table A.9.

To guard against erroneous tagging of non-AI and non-tech program descriptions, we matched keywords using case-insensitive regular expressions with word boundary detection, and accounted for exclusion patterns of common keywords that may appear in non-technical programs. For instance, the term "coding" which appears in some electronic health record administration programs that train on medical billing codes would not be tagged as having general technology content unless if it also contains a strong technology component (e.g. database management, software usage, systems administration). Of the 77,085 training programs, we flagged about 0.4% as having AI-specific content, 24% as having general technology content without AI, and 75% as having neither.

[20] These tags are not mutually exclusive.

We validate the results of our AI and technology classifications by linking the SOC occupation(s) associated with each training program to measures of AI exposure based on Eloundou et al. (2024) and Handa et al. (2025) ($AI^{\beta}_{Eloundou}$ and $AI_{Anthropic}$ respectively) as well as task-based automation indices from the Acemoglu and Autor (2011) task framework. We convert each measure to percentiles in order to compare the relative intensities across indices.

We show that, in line with expectations, AI-specific training programs are associated with target occupations that have higher AI-exposure, and are also associated with high degrees of routine and non-routine cognitive tasks. General technology training programs (excluding those that also contain AI content) are also associated with above median AI-exposure in their target occupations, although with lower magnitudes compared to those from AI-specific trainings. In contrast, other non-tech, non-AI trainings have are associated with target occupations with lower levels of AI-exposure, hovering around the median levels of the overall distribution.

Table A.8: Occupational Task Content of WIOA ETPL Programs

	Mean Percentile of Associated Occupations		
	AI Trainings	**Tech Trainings**	**Other Trainings**
A. AI Exposure Measures			
$AI^{\beta}_{Eloundou}$	83.9	75.6	50.6
$AI_{Anthropic}$	71.0	63.3	51.2
B. Automation Task Framework			
Routine Cognitive	60.4	68.2	55.2
Nonroutine Cognitive	69.6	57.4	56.4
Routine Manual	39.7	50.1	48.8
Nonroutine Manual	37.9	48.7	50.1

Notes: AI-exposure and automation task indices are reported as percentiles in terms of the overall distribution across all available 6-digit SOC codes. Higher percentiles pertain to higher levels of AI exposure or higher compatibility with each of the automation framework indices. $AI^{\beta}_{Eloundou}$ refers to SOC-level potential AI-exposure measures as described in Eloundou et al. (2024), while $AI_{Anthropic}$ refers to SOC-level realized AI exposure measures based on Anthropic's Economic Index Handa et al. (2025). Automation Task Framework Indices are based on SOC-level aggregations of the task-level measures described in Acemoglu and Autor (2011).

Table A.9: WIOA ETPL Program Tagging Keywords

Category	Keywords and Patterns
Panel A: Artificial Intelligence Keywords	
Core AI Terms	artificial intelligence, ai, machine learning, deep learning, neural network, neural net
AI Techniques	natural language processing, nlp, computer vision, image recognition, predictive analytics, predictive modeling, reinforcement learning
AI Applications	chatgpt, gpt, llm, large language model, generative ai, ai automation
Robotics	autonomous robot, intelligent robot, robot ai, autonomous system, unmanned system
Panel B: General Technology Keywords	
Programming & Software	programming, coding, software development, software engineering, web development, mobile development, app development, full stack, front-end, back-end
Programming Languages	python, java, javascript, c++, c#, ruby, php, sql, r, html5, css, react, angular, vue
Computer Science & IT	computer science, computer systems, information technology, it support, help desk, systems administration, systems analyst, technical support, it infrastructure
Cybersecurity & Networks	cybersecurity, information security, network security, security+, cisco, computer network, network administration, tcp/ip, firewall, cryptography, penetration test
Data & Analytics	data science, data analytics, data analysis, big data, data mining, data visualization, business intelligence, data engineer
Cloud & DevOps	cloud computing, aws, azure, google cloud, gcp, devops, docker, kubernetes
Databases	database administrator, database management, mongodb, postgresql, mysql, nosql, sql server
Design & Engineering	cad, cam, cad-cam, autocad, catia, solidworks, computer-aided design, engineering software
Medical Technology	ehr, electronic health record, medical coding software, healthcare software, health information
Industrial Automation	plc programming, programmable logic controller, cnc programming, computer control, industrial automation, scada
IT Certifications	comptia, a+ certification, network+, microsoft certified, cisco certified, ccna
Operating Systems	linux, unix, windows server

Appendix D. Quasi-Random Assignment to WIOA Caseworkers

Caseworkers are the gatekeepers of training under WIOA. Frontline WIOA staff—typically titled "case managers," "caseworkers," "career planners," or "employment specialists"—serve as the primary gatekeepers for training services under WIOA. When a worker enters an American Job Center (AJC), they first receive an initial assessment as required under 20 CFR § 678.430(a)(2). Training may be approved only if these staff determine, after an interview, evaluation, or assessment and career planning, that the participant meets the federal statutory criteria in 20 CFR § 680.210(a):

1. career services alone are insufficient for the worker to obtain self-sufficient employment;
2. training is needed to achieve the employment goal;
3. the worker has the ability to participate successfully in training.

These criteria are evaluated subjectively by each frontline staff member. Through Individual Employment Plans (IEPs) under 20 CFR § 680.170, check-ins at regular intervals provide multiple opportunities for frontline staff to recommend, reassess, or newly assign the participant to training as circumstances evolve. Otherwise, workers instead receive Wagner–Peyser basic career services, such as job search assistance, résumé preparation, and labor-exchange services.

Caseworkers differ substantially in how often they assign training. Local workforce development areas (LWDAs) layer additional requirements on to federal standards, such as in-demand occupation lists, local self-sufficiency wage thresholds, cost or duration caps on Individual Training Accounts (ITAs) that can be applied toward training, prerequisite or suitability requirements, and supervisory or committee approval—giving caseworkers considerable discretion over which workers qualify for training. This discretion produces large observable differences across administrative units: in the "WIA Gold Standard Evaluation" (Fortson et al., 2017), a randomized trial conducted across 28 randomly selected Local Workforce Investment Areas in 19 states, the share of eligible customers receiving staff-approved training ranged from 3 percent to 86 percent. Comparable patterns arise in the Trade Adjustment Assistance (TAA) program, where Hyman (2018) shows that variation in petition adjudication meaningfully affects which displaced workers ultimately receive TAA-funded training. Together, these findings demonstrate that frontline staff exercise substantial influence in determining access to publicly funded job training.

Assignment to caseworkers is effectively random due to operational constraints. The particular caseworker a jobseeker meets is typically determined by operational factors rather than participant characteristics, creating quasi-random variation in who conducts the assessment and training determination. AJCs generally assign walk-ins and scheduled appointments to the next available staff member based on real-time availability, rotating intake rosters, daily staffing matrices, or automated scheduling systems. These procedures are documented explicitly in local operating manuals. The San Bernardino County AJCC Desk Manual states that "every Workforce Development Specialist will be assigned new

cases on a rotation basis" and that customers are served by the "next WDS available" (Section 4, pp. 2–5). The Chicago–Cook Workforce Partnership Case Assignment Re-Assignment Instructions direct providers to assign customers using a rotation designed to "balance caseloads across staff" and to ensure assignment to "the first available career planner" (pp. 1–2). The San Francisco OEWD Workforce Services Protocols describe an intake flow in which jobseekers scheduling online or entering by walk-in are automatically routed to "the first open appointment slot" without regard to staff identity (Workforce Development Plan, 2017–2022, pp. 23–29). Because customers rarely request specific staff and AJCs prioritize rapid service delivery, these operational practices create a plausibly exogenous assignment mechanism that determines which caseworker applies the training-eligibility criteria specified under 20 CFR § 680.210(a).

Appendix D References

Chicago–Cook Workforce Partnership, *Case Assignment & Re-Assignment Instructions* Chicago–Cook Workforce Partnership. Internal Operating Guidance, pp. 1–2.

Fortson, Kenneth, David H. Greenberg, Annalisa Mastri et al., "Providing Public Workforce Services to Job Seekers: 30-Month Impact Findings from the WIA Gold Standard Evaluation," Report, Mathematica Policy Research 2017.

Hyman, Benjamin, "Can Displaced Labor Be Re-employed? Evidence from Quasi-Random Assignment to Trade Adjustment Assistance," 2018. Working paper.

San Bernardino County Workforce Development Department, *AJCC Desk Manual* San Bernardino County Workforce Development Department. Section 4, pp. 2–5.

San Francisco Office of Economic and Workforce Development, *Workforce Development Plan, 2017–2022* San Francisco Office of Economic and Workforce Development. Workforce Services Protocols, pp. 23–29.

U.S. Department of Labor, "Title 20, Code of Federal Regulations," Code of Federal Regulations. Sections 678.430, 680.170, and 680.210.

Appendix E. Caseworker Variation and Selection on Unobservables

In this appendix, we show that quasi-random assignment to unobserved caseworkers with heterogeneous assignment leniencies attenuates selection on unobservables, thereby strengthening the credibility of matching on observables.

Access to training is mediated by caseworker discretion in interpreting admission rules. Let individuals i be (as-good-as-randomly) assigned to caseworkers $j(i)$ within offices and time periods. Caseworkers differ in *leniency* toward allowing individuals into training programs. We model the decision to assign worker i to training (D_i) using a latent threshold rule:

$$D_i \;=\; \mathbf{1}\Big(\sigma_i + U_i + Z_{j(i)} > \xi_{j(i)}\Big) \tag{A.1}$$

where:

- σ_i: determinants of training assignment observed by both the caseworker and the econometrician (used in matching),
- U_i: determinants of training assignment and potential outcomes (Y(0) and Y(1)) unobserved by the econometrician, but potentially observed by the caseworker
- $Z_{j(i)}$: caseworker-specific leniency component (unobserved to the econometrician),
- $\xi_{j(i)}$: caseworker-specific assignment threshold reflecting heterogeneity in interpreting training approval rules (see Appendix D).

We assume that, conditional on observed characteristics (used in matching),

$$Z_{j(i)} \;\perp\; (U_i, Y_i(0), Y_i(1)) \;\Big|\; \sigma_i \tag{A.2}$$

reflecting quasi-random assignment of workers to caseworkers. Under decision rule (A.1), treatment depends on both U_i and the independent component $Z_{j(i)}$. Conditional on σ_i,

$$D_i = 1 \quad \Leftrightarrow \quad U_i > \xi_{j(i)} - \sigma_i - Z_{j(i)}.$$

Because $Z_{j(i)}$ is independent of U_i, assignment contains an exogenous component. Intuitively, if $Z_{j(i)}$ had zero variance, assignment would be a deterministic function of U_i and σ_i, producing strong selection on U_i. As the variance of $Z_{j(i)}$ increases, the role of U_i in determining D_i weakens. Formally, parameterize the dispersion of caseworker leniency with a scale parameter ($Z_{j(i)} = \lambda \tilde{Z}_{j(i)}$) where $\tilde{Z}$ has mean zero and unit variance, so that $\mathrm{Var}(Z) = \lambda^2$. Then

$$\frac{\partial}{\partial \lambda} \underbrace{\Big(\mathbb{E}[U_i \mid D_i = 1, \sigma_i] - \mathbb{E}[U_i \mid D_i = 0, \sigma_i]\Big)}_{\text{Selection on unobservables}} \;<\; 0 \tag{A.3}$$

so greater quasi-random caseworker variation dampens selection on unobservables, easing the identification burden placed on the matching variables σ_i.

Appendix F. Differencing Out Common Matching Bias

Many of the main findings in the paper are statements about heterogeneous effects. For example, we compare estimates for high AI-exposed training participants versus their nearest-neighbor (NN) control unit pairs, against low AI-exposed participants and their control unit pairs. While Figure 3 showed no pretrends in earnings for all subgroups relative to their nearest-neighbor control pairs, here we formalize how comparisons across groups allow us to further difference out any lingering imbalances, similar in spirit to a triple difference design.

Potential Outcomes Setup and Notation

Let i index individuals. Let $H_i \in \{0,1\}$ denote AI exposure group membership for individual i (e.g., $H{=}1 \equiv$ "High", $H{=}0 \equiv$ "Low"). Let $D_i \in \{0,1\}$ denote treatment (non-random). Potential outcomes are $Y_i(d,h)$ for $d \in \{0,1\}$, $h \in \{0,1\}$. The observed outcome is

$$Y_i \;=\; D_i\, Y_i(1,H_i) \;+\; (1{-}D_i)\, Y_i(0,H_i) \tag{A.4}$$

We want group-specific causal ATE parameters:

$$\tau_h \;\equiv\; \mathbb{E}[Y_i(1,h) - Y_i(0,h) \mid H_i = h]\,, \qquad h \in \{0,1\} \tag{A.5}$$

We estimate within-group "one-way" differences (treated – NN control within h):

$$\delta_h \;\equiv\; \mathbb{E}[Y_i(1,h) \mid D_i = 1] \;-\; \mathbb{E}[Y_i(0,h) \mid D_i = 0] \tag{A.6}$$

Lingering Selection Bias within Each Group

Add and subtract $\mathbb{E}[Y_i(0,h) \mid D_i = 1]$ to get the standard decomposition:

$$\delta_h \;=\; \underbrace{\tau_h}_{\text{causal ATT}} \;+\; \underbrace{\Big(\mathbb{E}[Y_i(1,h) \mid D_i = 1] - \mathbb{E}[Y_i(0,h) \mid D_i = 0]\Big)}_{\equiv\, B_h \text{ (lingering selection bias within } h)} \tag{A.7}$$

Each within-group estimate equals the true effect τ_h plus a group-specific selection bias B_h that remains even after nearest-neighbor matching.

Differencing Out Common Bias from Nearest-Neighbor Estimation

Consider the difference across the two groups of the within-group estimates:

$$\Delta \;\equiv\; \delta_1 - \delta_0 \;=\; \underbrace{(\tau_1 - \tau_0)}_{\text{heterogeneous effect of interest}} \;+\; \underbrace{(B_1 - B_0)}_{\text{residual/differential bias}} \tag{A.8}$$

The simple insight is that any component of the selection bias that is *common across groups* cancels when the focus is on differences in estimates. That is, if

$$B_h \;=\; \bar{B} + \varepsilon_h \tag{A.9}$$

with component $\bar{B}$ the matching bias shared by both groups h, and ε_h denotes group-specific selection effects across groups, then the bias becomes

$$\Delta - (\tau_1 - \tau_0) \;=\; (B_1 - B_0) \;=\; \varepsilon_1 - \varepsilon_0 \tag{A.10}$$

which in principle could be much smaller than the *level* bias $B_h = \bar{B} + \varepsilon_h$ that contaminates δ_h. Thus, the heterogeneous comparisons for Δ may be less biased than either one-way estimate δ_h. Formally, if the following assumption holds:

Assumption (Common selection on unobservables across H).

$$\mathbb{E}[Y_i(0,h) \mid D_i = 1, H_i = h] - \mathbb{E}[Y_i(0,h) \mid D_i = 0, H_i = h] \;=\; \bar{B} \quad \text{for } h \in \{0,1\} \tag{A.11}$$

then Δ may yield a less biased estimate for $\tau_1 - \tau_0$, and comparisons between groups yield improvements in bias over one-way differences.

Appendix G. Matching Estimator Details

We estimate treatment effects using the semiparametric nearest-neighbor matching estimator developed by Abadie and Imbens (2002). For each treated trainee i, the estimand is the individual treatment effect $\theta_i = Y_i(1) - Y_i(0)$, but only one potential outcome is observed. The estimator imputes the missing counterfactual outcome using outcomes of observationally similar control units.

Similarity between treated and potential control units is determined using a distance metric defined over pre-treatment covariates X. For each treated unit i, we search among control units j for the closest match in covariate space. To place covariates on a common scale and account for correlation among them, we use the Mahalanobis distance:

$$||X_i - X_j||_{\text{Mahalanobis}} = \sqrt{(X_i - X_j)' \Sigma_X^{-1} (X_i - X_j)}. \tag{A.12}$$

Matching is implemented using the algorithm described in Abadie et al. (2004), which pairs each treated observation with its nearest control neighbor according to this metric. This procedure uses a caliper (maximum allowable distance) to restrict estimation to regions of common *covariate support*: treated units for which no sufficiently close control match exists are excluded. Importantly, this form of overlap enforcement is based on distance in X rather than trimming based on estimated treatment probabilities, as in propensity-score approaches. In practice, fewer than 1% of treated observations are dropped due to lack of suitable matches, reflecting the large pool of potential comparison units drawn from Wagner–Peyser job search participants.

Alternative Matching and Weighting Approaches

While our baseline estimator follows the Abadie–Imbens nearest-neighbor framework, several alternative approaches are available and serve as robustness checks. First, estimation can be weighted inversely by match distance, placing greater weight on closer neighbors. This can improve covariate balance while preserving the geometry-based matching structure.

Second, matching can be expanded to include richer pre-treatment earnings histories. Prior earnings trajectories contain substantial information about unobserved heterogeneity and have been shown to absorb much of the selection bias in workforce program evaluations (Heckman et al., 1998; Lechner et al., 2011). Incorporating lagged earnings directly into the covariate vector X therefore provides a theory-guided expansion of the conditioning set.

Third, propensity-score approaches offer an alternative way to enforce overlap. Instead of matching on covariate distance, one can match or weight observations using estimated treatment probabilities $P(D = 1 \mid X)$. This approach enforces overlap in terms of treatment probabilities rather than geometric proximity. However, it requires estimating a first-stage model and can be sensitive when estimated propensities approach zero or one, where overlap violations occur. For example, Andersson et al. (2024) adopt a hybrid approach that combines both methods by performing nearest-neighbor matching on the propensity score itself. This approach leverages the dimension reduction of the propensity score while retaining the nearest-neighbor structure.)