

NBER WORKING PAPER SERIES

HOW RETRAINABLE ARE AI-EXPOSED WORKERS?

Benjamin G. Hyman
Benjamin Lahey
Karen Ni
Laura Pilossoph

Working Paper 34174
<http://www.nber.org/papers/w34174>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2025

We thank Joe Altonji, David Card, Lisa Kahn, Fabian Lange, Michael Lee, Kyle Myers, Steve Raphael, Daniel Rock, Wilbert van der Klauw, Till von Wachter, and seminar participants at the annual SOLE conference, Federal Reserve Banks of San Francisco and New York, 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. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of New York, the Federal Reserve System, or the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at <http://www.nber.org/papers/w34174>

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Benjamin G. Hyman, Benjamin Lahey, Karen Ni, and Laura Pilossoph. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Retractable are AI-Exposed Workers?

Benjamin G. Hyman, Benjamin Lahey, Karen Ni, and Laura Pilossoph

NBER Working Paper No. 34174

August 2025

JEL No. E0, E2, J6

ABSTRACT

We document the extent to which workers in AI-exposed occupations can successfully retrain for AI-intensive work. We assemble a new workforce development dataset spanning over 1.6 million job training participation spells from all US Workforce Investment and Opportunity Act programs from 2012–2023 linked with occupational measures of AI exposure. Using earnings records observed before and after training, we compare high AI exposure trainees to a matched sample of similar workers who only received job search assistance. We find that the average earnings return to training among AI-exposed workers is high, around \$1,470 per quarter. Low-exposure trainees capture higher returns, and trainees who target AI-intensive work face a 29% earnings return penalty relative to their high exposure peers who pursue more general training. We estimate that between 25% to 40% of occupations are “AI retrainable” as measured by its workers receiving higher pay for moving to more AI-intensive occupations—a large magnitude given the relatively low-income sample of displaced workers. Positive earnings returns in all groups are driven by the most recent years when labor markets were tightest, suggesting training programs may have stronger signal value when firms reach deeper into the skill market.

Benjamin G. Hyman
Federal Reserve Bank of New York
ben.hyman@ny.frb.org

Benjamin Lahey
New York University
bpl9631@nyu.edu

Karen Ni
Harvard University
Harvard Kennedy School
kni@g.harvard.edu

Laura Pilossoph
Duke University
and NBER
pilossoph@gmail.com

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 short paper evaluates the effectiveness of the US’s flagship workforce development program—the Workforce Innovation and Opportunity Act (WIOA), formerly the Workforce Investment Act (WIA)—in helping workers transition out of jobs facing AI-related pressure and into jobs with higher AI complementarity. We assemble a novel 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 compare with periods in which precursors to generative AI were 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. By merging in seminal measures of occupational AI exposure from [Brynjolfsson et al. \(2018\)](#) and their extensions to large language models (LLMs) and generative AI from [Eloundou et al. \(2024\)](#), we analyze the joint distribution of transitions in AI skill space and their associated earnings returns.

We find that earnings returns to training for AI-exposed workers are large and positive, but workers capture higher returns when they avoid targeting AI-intensive occupations in their next jobs. In an ideal experiment, these returns would be identified by random assignment. Absent random assignment, we follow the approach of [Rothstein et al. \(2022\)](#), matching each WIOA/WIA trainee to a control worker who only received job search assistance. The credibility of this design is strengthened by the fact that, although all participants receive job search assistance, entry into training is mediated by caseworkers, introducing an element of quasi-random assignment. In our context, the control group comprises nearest-neighbor matched workers from occupations with similar AI exposure in the same year. This addresses the potential “Ashenfelter dip” concern that a decline in earnings prior to training could upwardly bias estimates of post-training earnings gains. The nearest neighbor approach also allows us to separately examine

¹In 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.

outcomes for workers displaced from high-AI exposure occupations, and for those who are targeting transitions into AI-intensive occupations after training. That is, we can separately estimate the effects of more general (non-AI) and specific (AI-deepening) skill investments for workers who were highly exposed to AI prior to participating in training. Though not without limitations, the approach sheds early light on the understudied intersection of AI and job training.

We have four main findings. (i) The modal participant displaces from *top quintile* AI exposure occupations despite being relatively low income. This is not driven by the composition of the sample—training participants have composition similar to the CPS Unemployed and Displaced Worker samples. (ii) The mean earnings returns for training participants displaced from above-median (“high”) AI exposure occupations are large and positive—around \$1,470 per quarter relative to the control group and only about 25% lower than the returns for low AI exposure workers. (iii) Participants who target high AI exposure occupations for their next jobs also see significant returns, although with a penalty relative to those who pursue more general skills. Returns for participants deepening AI-specific skills remain high, at roughly \$1,040 per quarter, but are approximately 29% lower compared to the average returns for high AI exposure workers who pursue more general skills. (iv) Positive returns in both groups are largely driven by the most recent years when labor markets were tightest, suggesting training programs may have stronger signal value when firms are forced to reach deeper into the skill market.

Using the matched sample, we also construct an AI Retraining Index (AIR) which ranks occupations by the share of workers who can 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 between 25% to 40% 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 main finding that workers are more successful when they avoid targeting AI occupations for their next job contributes 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.² They are also silent on the question of whether workers would be better off avoiding AI occupations altogether. Firm-level studies are also naturally limited from drawing inferences across a wider set of workers and firms. 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 is an important driver of positive returns also contributes to the literature on active labor market policy evaluations, especially those 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. [Brynjolfsson et al. \(2018\)](#), [Eloundou et al. \(2024\)](#)) have a natural mapping to task content measures from the preceding automation literature ([Acemoglu and Autor \(2011\)](#)). Specifically, we show that AI exposure is correlated with routine cognitive tasks while generative AI and LLM exposure is correlated with both routine and non-routine cognitive task measures. 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).

2 Data and Measurement

Determining whether individuals can be retrained for AI-related work is challenging because we typically lack administrative data on how AI is used by workers within firms.³ We make progress on this challenge by developing a dataset that can longitudinally track transitions of

²In a prominent example, [Brynjolfsson et al. \(2025\)](#) finds that customer service ChatBots helped less experienced workers catch up to experienced colleagues.

³While 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.

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. 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.

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. 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

⁴Popular 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.

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 case workers to enroll in Title I training services (29 U.S.C. §3174(c)(3)(A)(i)).⁵ To account for potential 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 ES-202 administrative records 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.⁶ One limitation of the PIRL dataset is that the 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, our final sample contains 109,038 total training spells, of which 108,215 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

We use two seminal occupational measures of AI exposure that we merge to WIOA/WIA participant occupations to study AI transitions. The first is [Brynjolfsson et al. \(2018\)](#)’s “Suitability for Machine Learning” (SML) measure which is independently calculated for each pair of tasks (2,069 O*NET “work activities”) and occupations (964 6-digit SOC occupations). 7 independent human annotators evaluate each task-occupation pair across 23 Likert-scale questions that measure the likelihood that a machine could execute the task

⁵Because the case worker is unobserved, we cannot use an examiner design to estimate effects as in [Humlum et al. \(2023\)](#).

⁶All 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.

in a given occupation, where a score of 5 reflects high suitability for machine learning and 1 low suitability. The SML measure uses O*NET “importance weights” that determine the relevance of each task for a given occupation. We aggregate SML importance-weighted scores by occupation by taking the mean across all questions and annotators, resulting in a single time-invariant SML score for each 6-digit occupation. Among the unique occupations spanned by the WIOA/WIA dataset (prior to program participation), the mean and median SML scores are 3.47 and 3.46 with a standard deviation of 0.10.

While the SML score is a pioneering approach to AI exposure, it was developed in 2017. To capture whether subsequent advances in large language models (LLMs) including generative technology such as ChatGPT can execute the same task-occupation pairs, we draw from a more recent measure developed by Eloundou et al. (2024) which we denote SML_{β}^{LLM} . Eloundou et al. (2024) expand on the earlier SML methodology by asking whether generative AI and LLMs specifically can reduce the time required to complete a given task. Both annotators and ChatGPT itself evaluate the likelihood that (a) an LLM would reduce the time taken to perform the task by at least 50% and (b) additional software could be developed on top of the LLM to reduce the time taken by at least 50%.⁷

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 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. However, one novel contribution of our paper is that we show that the two AI exposure measures used in this paper (SML and SML_{β}^{LLM}) correlate with foundational measures used in the Acemoglu and Autor (2011) task framework. In Appendix Figure A.1, we illustrate this link by plotting binscatter plots which reveal that SML is strongly correlated with *routine-cognitive* tasks, while SML_{β}^{LLM} is correlated with both *routine-* and *non-routine cognitive* tasks. That is, we formally document that AI tasks are cognitive tasks. This observation is useful 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

⁷We use the “ β ” weighting scheme which calculates a variable $SML_{\beta}^{LLM} \in [0, 1] = (a) + 0.5 * (b)$ that captures the share of an occupation’s tasks that can be done by LLMs. This weighting scheme is chosen because it has the desirable validation that SML_{β}^{LLM} is positively correlated with average occupational incomes.

years (more aligned with automation) and increasingly 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 by whether their pre-enrollment occupation was above or below the 6-digit median occupation-level SML score, resulting in a “high” AI exposure group and a “low” AI exposure group.⁸ [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 108,215 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 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.⁹

[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 [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.

⁸We 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.

⁹Occupations spanned by our full set of WIOA/WIA training participants account for around 93% of the CPS labor force on a monthly basis.

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. These plots reveal that after a mild “Ashenfelter dip” in the run up to training, post-training earnings are mildly higher in most years, with much greater values in more recent years (indicated with heavier lines) when labor markets were tighter. This pattern appears to be more dominant in driving earnings patterns than heterogeneity across low and high AI exposure occupations, though the data do point to slightly lower post-training earnings among highly exposed groups. When cutting the data by high AI intensity (exposure) in the target occupation of training—that is, the desired occupation after training—we see a similar pattern. Overall, the stronger patterns in years in which labor markets were tighter may suggest that training programs may have stronger signal value when firms are forced to reach deeper down in the skill market. We adjust for potential selection in these patterns in our nearest neighbor matching strategy discussed in [Section 3](#).

Job Flows Across AI Quintiles. In a final descriptive, we analyze the flows of job trainees across AI skill space in [Figure 2](#). On the x-axis, we split trainees in our sample into five quintiles of AI exposure prior to training participation. Before analyzing flows between quintiles, it is worth noting that the modal WIOA/WIA training participant is displaced from a top 20th percentile (i.e. top quintile) occupation 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 (see [Appendix Table A.3](#) for a detailed breakdown by AI exposure quintile). Further decomposition of the bars into destination occupation quintiles after training suggests that most workers are moving to weakly higher AI skills relative to the AI content of their prior jobs, with only a small share of workers remaining in their same 6-digit occupation before and after training. Overall, this pattern points to substantial mobility across occupations after training.¹⁰ 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 flows into more intensive AI work are consistent with successful AI retraining, WIOA/WIA workers may be positive selected, either by virtue of the Ashenfelter dip which depresses earnings prior to training participation, or due to changes in trainee composition over time that explain higher returns in more recent years. 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 credibility of this design is strengthened by the fact that, although all participants receive job search assistance, entry into training is mediated by caseworkers, introducing an element of quasi-random assignment.

Because our goal is to attain heterogeneous estimates across AI exposure groups, 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 B](#) that heterogeneous comparisons across groups help difference out lingering selection bias from the matching procedure in any given group.

We use the semiparametric 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 (see [Appendix C](#) for details). This has the useful property that violators of the “covariate overlap” assumption are dropped, resulting in treated units only being used

¹⁰In [Appendix Figure A.3](#), we show that while the share of workers flowing to higher AI occupations has remained stable, those making the transition have received growing earnings returns over time.

if the covariates share a sufficiently common covariate support. In practice, less than 1% of our final sample is dropped due to overlap, a high match success rate that arises due to the abundance of potential control units available in the Wagner-Peyser program.

We require certain variables to be an “exact match” for each nearest neighbor pair, and then rely on “fuzzy matching” for remaining covariates:

- *Exact Match Covariates*: calendar year of program exit (controls for demand environment), SML 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
- *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

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.5](#) 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 conducting 22 t-tests), differences in mean values for treated and control groups are, for the most part, economically 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 1 to 2 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.

4 Main Results

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 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. Panel (B) shows analogous effects, but for participants listing high AI intensity (exposure) occupations as their target occupation after job training, versus low AI intensity occupations, relative to their control group pairs.

In both panels, prior earnings differences are well balanced in all sub-groups despite the exclusion of lagged outcomes from our covariate match lists. This mitigates concerns about selection on unobserved heterogeneity, and is consistent with caseworkers adding an element of quasi-random gatekeeping into training. Starting with Panel (A), estimated earnings returns for high initial AI exposure training participants are slightly muted 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 about \$1,450-\$1,650 higher relative to their matched control group pairs, or on average \$6,225 annually on a base of roughly \$44,000 prior to participation (see [Table A.1](#) for baseline real wages). That is, among high AI exposure workers, those who train have about 14% 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 concentrated in the most recent years. [Appendix Table A.6](#) shows pooled regression estimates associated with these earnings returns before and after participation.

Training participants in Panel (A) are free to choose any occupation after training, including general skill training that facilitates avoiding AI occupations that may be at risk of future displacement. By contrast, Panel (B) shows estimates based on whether the target occupation (the individual’s desired occupation as a result of training) is high or low AI intensity. Here, we see a substantial penalty for workers that choose to retrain in high AI intensity occupations, thus deepening AI-specific human capital relative to those who appear to instead elect general training in Panel (A). Participants targeting high AI occupations in training programs in Panel (B) make on average \$1,039 more than their matched control group pairs, which is about a 29% lower return to training relative to high AI exposure workers that retrain in all occupations (see [Appendix Table A.6](#)). This pattern in relative returns is robust to using a wide range of percentile thresholds for defining “high exposure,” not just the median (see [Appendix Figure A.5](#)).

One implication of these results is that AI-exposed workers appear highly adaptable when supported by a training program, but may face frictions in adapting to high AI intensity work in particular. Indeed, the short-run returns to all four types of workers

featured in [Figure 3](#) are unambiguously large with respect to the literature (e.g. comparing to active labor market program treatment effect estimates considered in the meta-study by [Card et al. \(2018\)](#), who concluded they are “ineffective” in the short term). Since the effects are largely driven by the later years in our training sample in which market tightness was unprecedentedly high, our evidence suggests that training programs may interact in important ways with the state of the labor market. For example, the large returns when tightness is high are consistent with training programs containing a stronger signal value when firms reach deeper into the skill market. 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.

4.1 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 the final two columns of [Table A.7](#), we show the share of training spells for which both AI intensity and earnings are weakly greater, separately for the training participant group and the training group relative to matched nearest neighbors. We find that about 26% - 28% of workers are AI retrainable, with this range improving slightly to 28% - 34% when using the LLM measure developed by [Eloundou et al. \(2024\)](#). 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 a range varying from about 25% to 40%. These AI retrainable occupations cover 38% 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 AI Retraining (AIR) Index

To further unpack the mechanisms underlying the estimated large degree of adaptability to AI shocks through retraining, we develop an “AI Retraining 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.¹¹ The index is as follows

$$AIR_i = \sum_j \frac{1}{n_{ij}} \Delta SML_{ij}^\sigma * \mathcal{M}(\Delta \log(\overline{Y}_{ij})) \quad (1)$$

where n_{ij} is the number of trainees flowing from occupation i to j , capturing probability of reemployment in occupations of varying AI intensities. $\Delta SML_{ij}^\sigma = (SML_j^\sigma - SML_i^\sigma)_{p=1} - (SML_j^\sigma - SML_i^\sigma)_{p=0}$ is the difference in standardized mean SML scores associated with each trainee’s occupational transition from i to j relative to the change in SML for their matched control group unit. $\Delta \log(\overline{Y}_{ij}) = \frac{1}{p_{ij}} \sum_p \Delta \log(Y_{p=1}^{i \rightarrow j}) - \Delta \log(Y_{p=0}^i)$ 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 SML changes by converting the most negative earnings returns into low positive values:

$$\mathcal{M}(\Delta \log(\overline{Y}_{ij})) = \frac{\Delta \log(\overline{Y}_{ij}) - \min(\Delta \log(\overline{Y}_{ij}))}{\max(\Delta \log(\overline{Y}_{ij})) - \min(\Delta \log(\overline{Y}_{ij}))} \quad (2)$$

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. SML) 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.¹²

To see this clearly, [Figure 4](#) presents the overall 2-digit rankings of occupations by

¹¹In 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.

¹²In 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.

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 *Legal, Computer and Mathematics*, as well as *Art, Design, and Entertainment/Sports Media* professions. These also rank highly in Anthropic’s Economic Index measuring the task use of AI by occupation (Handa et al. (2025)). 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, rather than from left to right in earnings returns. In a final exercise, [Appendix Figure A.6](#) shows decompositions for the modal 6-digit occupation within three illustrative occupations, and demonstrates that occupations vary tremendously in their adaptability to AI, with some workers deepening their AI skills and others avoiding them.

6 Discussion

We document 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 that are only about 25% lower than the returns for low AI exposure workers. However, training participants who choose to deepen in AI-specific skills face a penalty for doing so, with 29% lower returns than those pursuing general training. This suggests that the predominant adjustment margin for AI-exposed workers is not through deepening their AI skills, but rather avoiding AI altogether.

Despite the inferred penalty from AI skill deepening, earnings returns to WIOA/WIA training are high for all AI exposure levels. We estimate that between 25% to 40% 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. Positive earnings returns in all groups are driven by the most recent years when labor markets were tightest, suggesting training programs may have stronger signal value when firms reach deeper into the skill market.

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. Future work will need to expand the analysis to consider whether workers who receive on-the-job training, as well as workers more aligned with white-collar jobs, remain resilient as AI tools continue to diffuse throughout the economy.

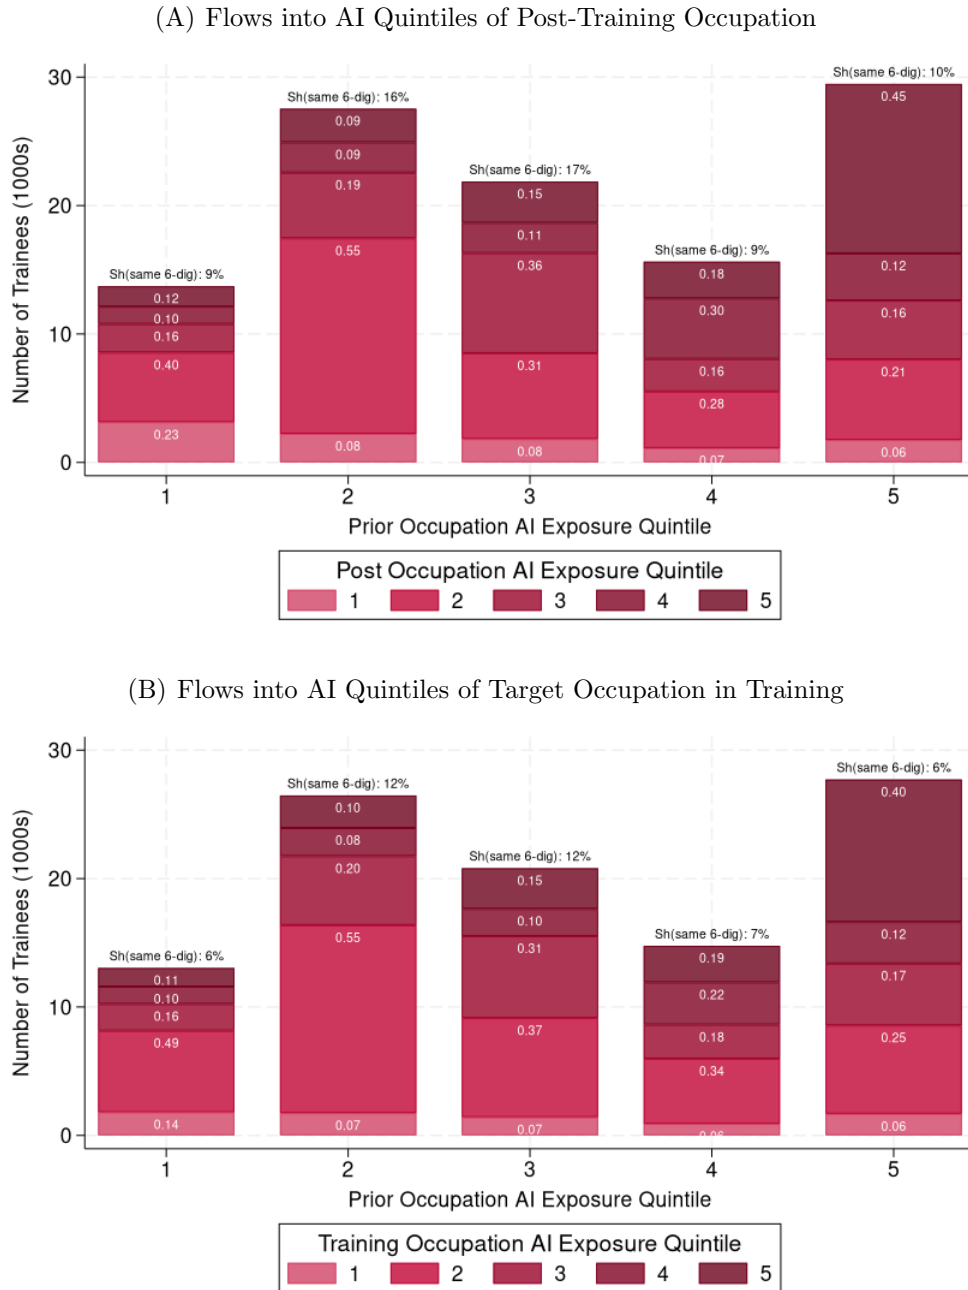
Figures and Tables

Figure 1: Earnings Selection Patterns for WIOA/WIA Trainees by Program Year



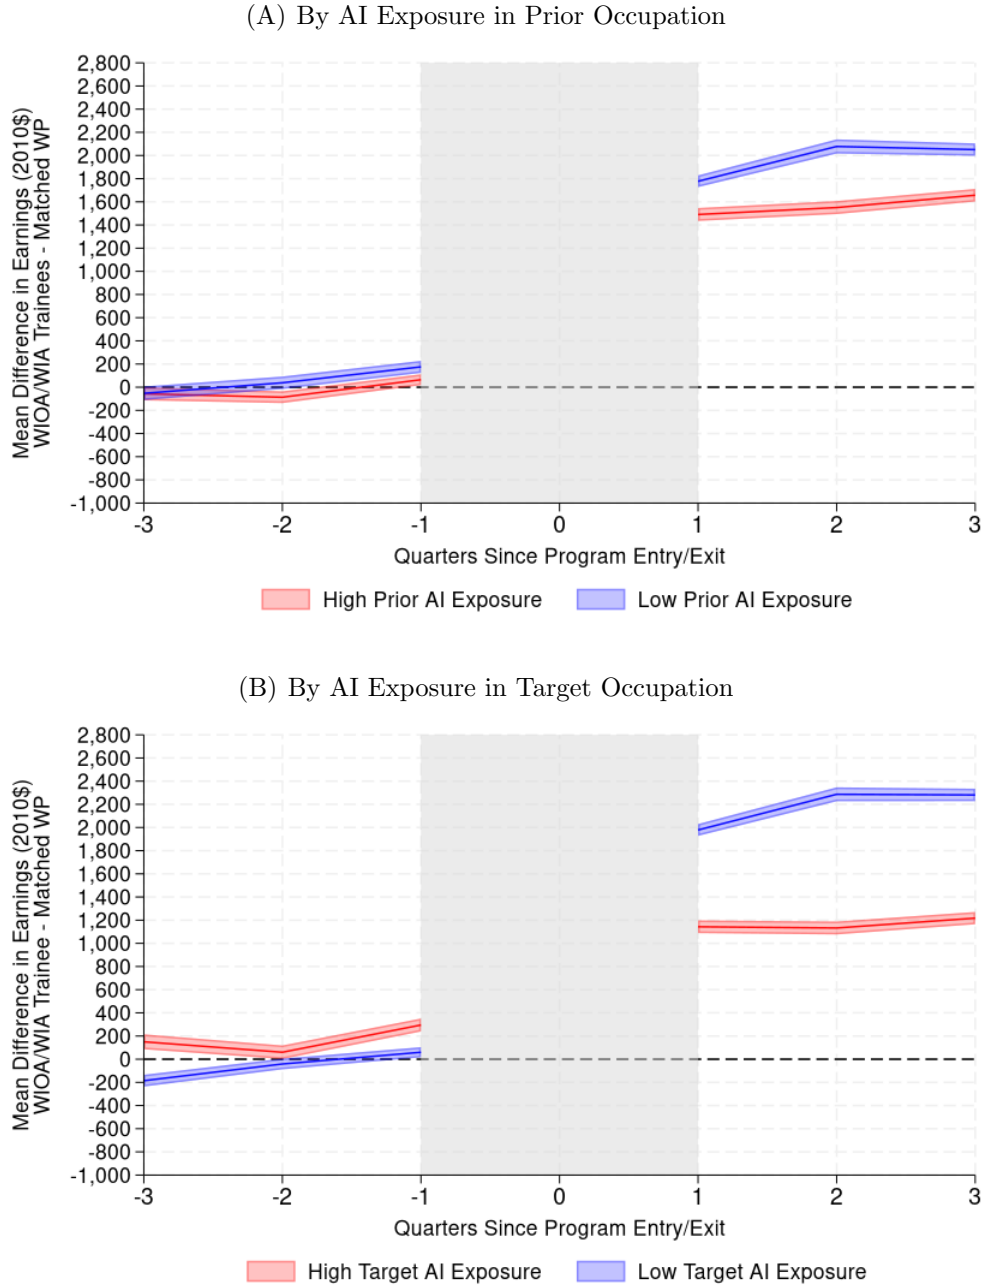
Notes: Figures plot mean quarterly earnings (CPI-deflated to 2010 real dollars) for all WIOA/WIA training participants in the main analysis sample, subject to the sample restrictions discussed in [Section 2.3](#). High and Low AI exposure is measured as workers having occupations above or below the median 6-digit occupational AI exposure measure from [Brynjolfsson et al. \(2018\)](#) (see text for details). Panels (A) and (B) show results for occupations prior to training participation, while panels (C) and (D) show results for desired or “target” occupation in next job after training. Data are observed for three quarters before training spell entry (red series) and three quarters after training spell exit (blue series), but not in the intervening period (shaded gray). Heavier lines correspond to more recent years while fainter lines correspond to earlier years.

Figure 2: Flows Across AI-Exposed Occupations Before and After Job Training



Notes: Figure show transitions between occupations prior to training and occupations after training (Panel A) or desired/targeted occupations in training programs (see text for details). X-axis reflects quintiles of AI exposure in occupation prior to training, while darker red colors reflect higher AI quintiles in destination or target occupation. High and Low AI exposure is measured as workers having occupations above or below the median 6-digit occupational AI exposure measure from Brynjolfsson et al. (2018) (see text for details).

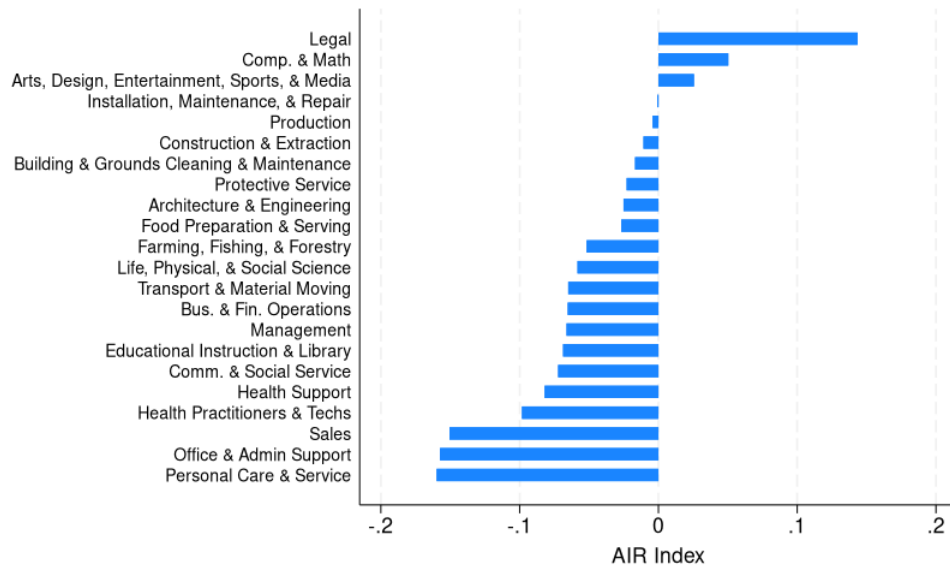
Figure 3: Nearest-Neighbor Matched Returns to WIOA/WIA: Pooled Across Years




Notes: Panel (A) shows estimates for the effects of WIOA/WIA training relative to the matched Wagner-Peyser control group, separately for high (red) and low (blue) AI exposure groups. Each series shows estimates that reflect mean differences between training and control group participants, weighted by the inverse Mahalanobis distance between training and control unit pairs, separately estimated in each quarter with t-tests along with equal-variance 95% confidence intervals. Panel (B) shows analogous effects, but for participants listing high AI intensity (exposure) occupations as their target occupation after job training, versus low AI intensity occupations.

Figure 4: AIR Index Rankings by 2-digit SOC Code Prior to Training

(A) Index Rankings



(B) Decomposition



Notes: Panel (A) ranks 2-digit SOC occupations based on the extent to which displaced workers from these occupations are subsequently retrained into higher AI-intensity occupations. The gradient scale in Panel (B) reflects higher rankings in darker red colors, and decomposes rankings based AI skill movements (y-axis) and earnings movements (x-axis). See text for details.

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.
- 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, 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.
- 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.
- 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.
- 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.

Online Appendix (Not for Publication)


Appendix A. Additional Results

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

(A) SML scores at 6-digit (Occupation Level)

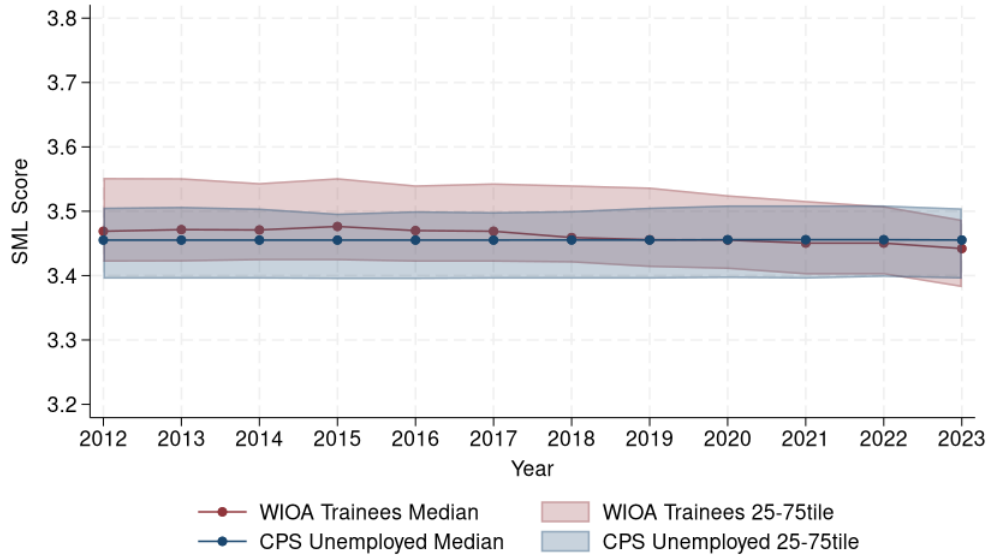


(B) SML-LLM values vs. AA Index




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—Suitability for Machine Learning (SML) at the 6-digit occupation level following Brynjolfsson et al. (2018). The y-axis for panel (B) represents the LLM extension of AI exposure which we call “SML-LLM” following Eloundou et al. (2024). Correlations between AA task shares and SML 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 (“Suitability for Machine Learning” from Brynjolfsson et al. (2018)) for WIOA/WIA trainees’ occupations prior to training, and CPS unemployed workers’ occupations prior to their current unemployment spell.

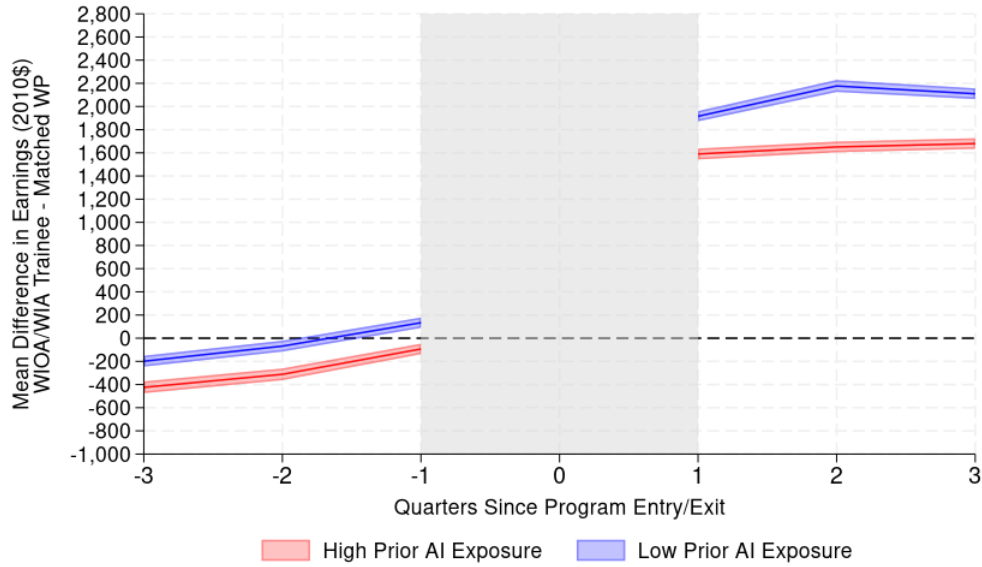
Figure A.3: Share of Positive SML and Earnings Changes over Time



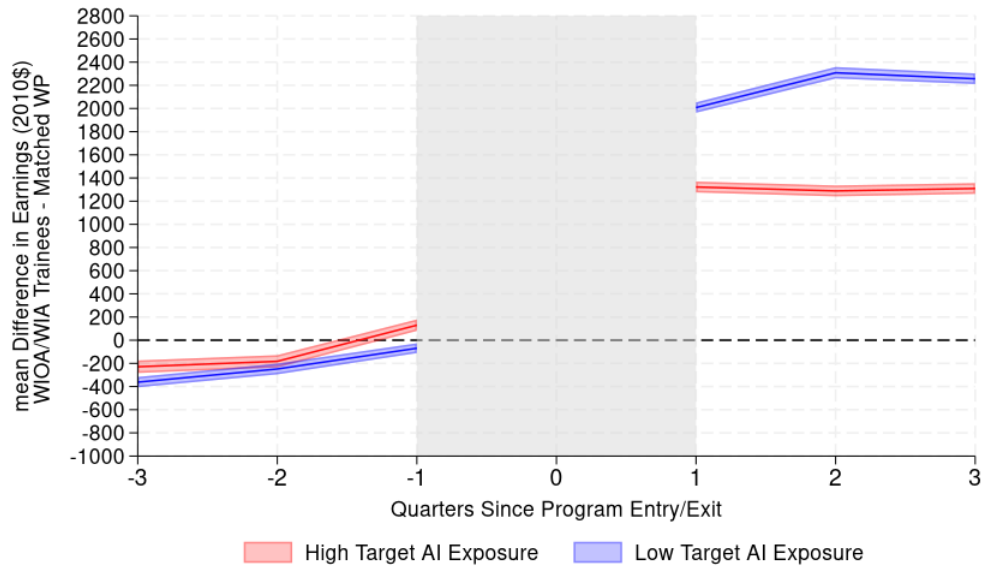
Notes: Figure plots share of WIOA/WIA training participants moving to higher AI exposure occupations (“Suitability for Machine Learning” scores from Brynjolfsson et al. (2018)) after training by calendar year of participation, and jointly receiving higher earnings from moving to higher AI exposure values.

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

(A) By AI Exposure in Prior Occupation



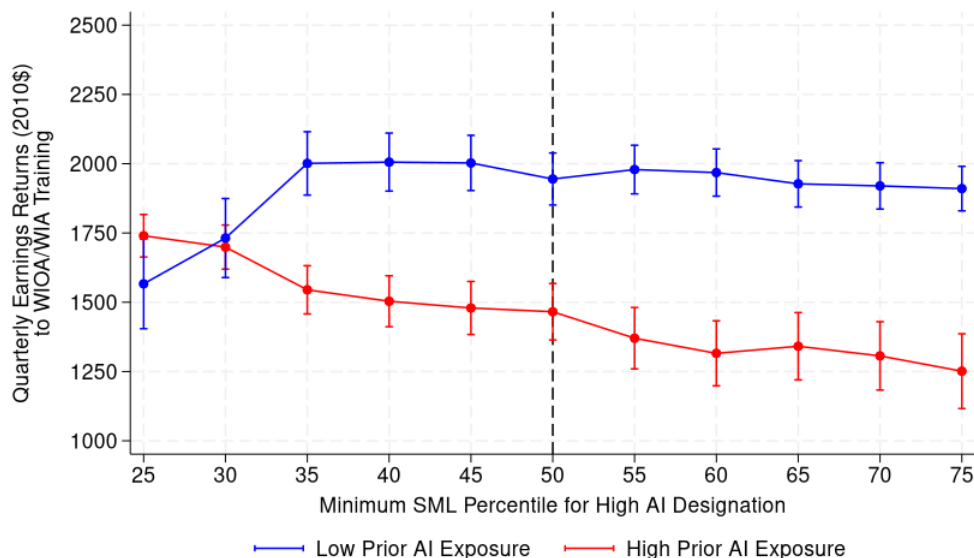
(B) By AI Exposure in Target Occupation




Notes: Panel (A) shows estimates for the effects of WIOA/WIA training relative to the matched Wagner-Peyser control group, separately for high (red) and low (blue) AI exposure groups, including earnings observations equal to zero. Each series shows estimates that reflect mean differences between training and control group participants, weighted by the inverse Mahalanobis distance between training and control unit pairs, separately estimated in each quarter with t-tests along with equal-variance 95% confidence intervals. Panel (B) shows symmetric effects, but for participants listing high AI intensity (exposure) occupations as their target occupation after job training, versus low AI intensity occupations.

Figure A.5: Robustness of Earnings Returns Estimates to Definition of High AI Exposure

(A) By AI Exposure in Prior Occupation




(B) By AI Exposure in Target Occupation




Notes: Figures show sensitivity of our nearest-neighbor matched estimates for the quarterly returns to training (pooled over the three quarters after program exit) to varying the cutoff for defining high AI exposure. Dashed vertical line corresponds to our preferred estimates using the median SML score to define AI exposure. Heteroskedasticity-robust 95% confidence intervals are shown in whiskers.

Figure A.6: 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 20/23)



(C) Customer Service Representatives (Office & Admin Support = 2-Digit AIR Rank 22/23)



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.

Appendix A.1 Decomposition of AIR Index

In [Appendix Figure A.6](#), we show 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 20/23) whose dominant 6-digit occupation is *Nursing Assistants*, and *Office and Admin Support* (AIR Rank 22/23) whose dominant occupation is *Customer Service Representatives*.

[Appendix Figure A.6](#) 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.

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 (1) | Δ /SE (2) | # Training Spells (3) | Mean/SD (4) | Δ /SE (5) | # Training Spells (6) | Δ /SE | % Diff |
| A. Demographic | | | | | | | | |
| Female | 0.57 [0.50] | | 53,232 | 0.44 [0.50] | | 54,983 | 0.13 (0.0030) | 22.1 |
| Age | 40.1 [11.3] | | 53,232 | 37.8 [10.7] | | 54,983 | 2.28 (0.067) | 5.68 |
| Asian | 0.040 [0.20] | | 53,232 | 0.027 [0.16] | | 54,983 | 0.012 (0.0011) | 30.8 |
| Black | 0.25 [0.43] | | 53,232 | 0.29 [0.45] | | 54,983 | -0.041 (0.0027) | -16.8 |
| White | 0.60 [0.49] | | 53,232 | 0.57 [0.49] | | 54,983 | 0.025 (0.0030) | 4.19 |
| Hispanic/Latino | 0.16 [0.37] | | 53,232 | 0.15 [0.36] | | 54,983 | 0.011 (0.0022) | 7.05 |
| Disability Status | 0.029 [0.17] | | 53,232 | 0.025 [0.16] | | 54,983 | 0.0037 (0.00098) | 12.9 |
| No HS Diploma/GED | 0.034 [0.18] | | 53,232 | 0.055 [0.23] | | 54,983 | -0.022 (0.0013) | -64.5 |
| HS Diploma/GED | 0.42 [0.49] | | 53,232 | 0.52 [0.50] | | 54,983 | -0.10 (0.0030) | -24.9 |
| Some College | 0.20 [0.40] | | 53,232 | 0.19 [0.39] | | 54,983 | 0.0053 (0.0024) | 2.71 |
| College Degree Plus | 0.35 [0.48] | | 53,232 | 0.23 [0.42] | | 54,983 | 0.12 (0.0027) | 34.3 |
| B. Social Benefits | | | | | | | | |
| Disadvantaged Adult Status | 0.48 [0.50] | | 53,232 | 0.58 [0.49] | | 54,983 | -0.10 (0.0030) | -21.2 |
| Dislocated Worker Status | 0.52 [0.50] | | 53,232 | 0.42 [0.49] | | 54,983 | 0.11 (0.0030) | 20.3 |
| Low-income Status | 0.42 [0.49] | | 53,232 | 0.46 [0.50] | | 54,983 | -0.041 (0.0030) | -9.77 |
| TANF Recipient | 0.0094 [0.096] | | 53,232 | 0.0092 [0.095] | | 54,983 | 0.00019 (0.00058) | 2.02 |
| SNAP Recipient | 0.21 [0.41] | | 37,768 | 0.23 [0.42] | | 40,865 | -0.020 (0.0030) | -9.25 |
| Other Public Assistance Recipient | 0.25 [0.43] | | 41,854 | 0.27 [0.44] | | 45,662 | -0.019 (0.0030) | -7.72 |
| C. Pre-Separation Employment | | | | | | | | |
| Employed at Participation | 0.30 [0.46] | | 53,232 | 0.37 [0.48] | | 54,983 | -0.069 (0.0029) | -23.4 |
| Real Earnings 1Q Before | 10687.3 [11522.8] | | 53,232 | 9287.5 [10355.7] | | 54,983 | 1399.8 (66.6) | 13.1 |
| Real Earnings 2Q Before | 11139.5 [9657.6] | | 53,232 | 9966.4 [9522.3] | | 54,983 | 1173.2 (58.3) | 10.5 |
| Real Earnings 3Q Before | 11020.0 [9689.5] | | 53,232 | 9785.8 [9028.0] | | 54,983 | 1234.2 (56.9) | 11.2 |
| Quarters in WIOA/WIA | 5.05 [3.43] | | 53,232 | 4.60 [3.23] | | 54,983 | 0.45 (0.020) | 8.96 |

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 108,215 participation spells.

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

| | Nonmissing Occ Codes | | | Missing Occ Codes | | | Nonmiss - Miss | |
|-------------------------------------|----------------------|---------------------|--------------------------|----------------------|---------------------|--------------------------|----------------------|--------|
| | Mean/SD (1) | Δ /SE (2) | # Training Spells (3) | Mean/SD (4) | Δ /SE (5) | # Training Spells (6) | Δ /SE | % Diff |
| A. Demographic | | | | | | | | |
| Female | 0.50 [0.50] | | 108,215 | 0.50 [0.50] | | 469,123 | -0.0010 (0.0017) | -0.20 |
| Age | 38.9 [11.1] | | 108,215 | 38.1 [11.0] | | 388,251 | 0.84 (0.038) | 2.17 |
| Asian | 0.042 [0.20] | | 86,669 | 0.035 [0.18] | | 403,080 | 0.0069 (0.00070) | 16.6 |
| Black | 0.32 [0.47] | | 89,940 | 0.31 [0.46] | | 407,087 | 0.0099 (0.0017) | 3.10 |
| White | 0.67 [0.47] | | 94,208 | 0.64 [0.48] | | 421,438 | 0.034 (0.0017) | 4.99 |
| Hispanic/Latino | 0.16 [0.37] | | 103,164 | 0.16 [0.37] | | 424,666 | 0.0056 (0.0013) | 3.40 |
| Disability Status | 0.027 [0.16] | | 105,648 | 0.027 [0.16] | | 443,843 | 0.00018 (0.00056) | 0.67 |
| No HS Diploma/GED | 0.045 [0.21] | | 108,215 | 0.058 [0.23] | | 469,123 | -0.014 (0.00077) | -30.4 |
| HS Diploma/GED | 0.47 [0.50] | | 108,215 | 0.46 [0.50] | | 469,123 | 0.0052 (0.0017) | 1.11 |
| Some College | 0.19 [0.39] | | 108,215 | 0.19 [0.39] | | 469,123 | 0.0028 (0.0013) | 1.46 |
| College Degree Plus | 0.29 [0.45] | | 108,215 | 0.28 [0.45] | | 469,123 | 0.0087 (0.0015) | 2.98 |
| B. Social Benefits | | | | | | | | |
| Disadvantaged Adult Status | 0.53 [0.50] | | 108,215 | 0.65 [0.48] | | 469,123 | -0.12 (0.0016) | -22.5 |
| Dislocated Worker Status | 0.47 [0.50] | | 108,215 | 0.38 [0.48] | | 469,123 | 0.090 (0.0016) | 19.3 |
| Low-income Status | 0.44 [0.50] | | 108,215 | 0.43 [0.49] | | 437,578 | 0.017 (0.0017) | 3.86 |
| TANF Recipient | 0.011 [0.10] | | 92,534 | 0.013 [0.11] | | 391,131 | -0.0025 (0.00041) | -23.5 |
| SNAP Recipient | 0.22 [0.42] | | 78,633 | 0.21 [0.41] | | 237,363 | 0.010 (0.0017) | 4.70 |
| Other Public Assistance Recipient | 0.26 [0.44] | | 87,516 | 0.22 [0.41] | | 371,110 | 0.042 (0.0016) | 16.1 |
| C. Pre-Separation Employment | | | | | | | | |
| Employed at Participation | 0.33 [0.47] | | 108,215 | 0.40 [0.49] | | 469,123 | -0.072 (0.0016) | -21.7 |
| Real Earnings 1Q Before | 9975.8 [10967.7] | | 108,215 | 9645.9 [13024.5] | | 469,123 | 329.9 (42.7) | 3.31 |
| Real Earnings 2Q Before | 10543.2 [9607.0] | | 108,215 | 10096.0 [11923.0] | | 469,123 | 447.3 (38.9) | 4.24 |
| Real Earnings 3Q Before | 10392.9 [9379.5] | | 108,215 | 9864.1 [10391.4] | | 469,123 | 528.7 (34.4) | 5.09 |
| Quarters in WIOA/WIA | 4.82 [3.34] | | 108,215 | 4.39 [3.15] | | 469,123 | 0.43 (0.011) | 8.96 |

Notes: Columns (1) and (2) show summary statistics for the main analysis sample of job trainees (N=108,215 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: Top Occupations in Each AI Exposure Quintile Ranked by Number of Trainees

| Rank | 6-Digit SOC | Occupation (No. of Trainees) |
|---|--------------------|--|
| <i>AI Exposure Quintile 1 (Low Exposure)</i> | | |
| 1 | 537062 | Laborers and freight, stock, and material movers, hand (1,923) |
| 2 | 537051 | Industrial truck and tractor operators (793) |
| 3 | 537064 | Packers and packagers, hand (616) |
| <i>AI Exposure Quintile 2</i> | | |
| 1 | 533032 | Heavy and tractor-trailer truck drivers (1,934) |
| 2 | 519198 | Helpers – production workers (1,487) |
| 3 | 292061 | Licensed practical and vocational nurses (1,468) |
| <i>AI Exposure Quintile 3</i> | | |
| 1 | 412031 | Retail sales (1,430) |
| 2 | 19199 | Managers, all other (1,326) |
| 3 | 533033 | Light truck drivers (1,276) |
| <i>AI Exposure Quintile 4</i> | | |
| 1 | 533031 | Drivers/sales workers (829) |
| 2 | 519061 | Inspectors, testers, sorters, samplers, weighers (748) |
| 3 | 514041 | Machinists (629) |
| <i>AI Exposure Quintile 5 (High Exposure)</i> | | |
| 1 | 434051 | Customer service representatives (3,851) |
| 2 | 412011 | Cashiers (1,993) |
| 3 | 439061 | Office clerks, general (1,379) |

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 parentheses. See text for details.

Table A.4: Balance Table for Matched Pairs with High AI Exposure Prior to Participation

| | Trainees | | | Matched Wagner Peyser | | | Trainees - Wagner Peyser | |
|-------------------------------------|----------------------|---------------------|-----------------|-----------------------|---------------------|-----------------|--------------------------|--------|
| | Mean/SD (1) | Δ /SE (2) | # Spells (3) | Mean/SD (4) | Δ /SE (5) | # Spells (6) | Δ /SE | % Diff |
| A. Demographic | | | | | | | | |
| Female | 0.57 [0.50] | | 53,232 | 0.59 [0.49] | | 53,232 | -0.025 (0.0030) | -4.41 |
| Age | 40.1 [11.3] | | 53,232 | 41.5 [11.1] | | 53,232 | -1.44 (0.069) | -3.60 |
| Asian | 0.040 [0.20] | | 53,232 | 0.034 [0.18] | | 53,232 | 0.0058 (0.0012) | 14.6 |
| Black | 0.25 [0.43] | | 53,232 | 0.25 [0.43] | | 53,232 | -0.0056 (0.0026) | -2.28 |
| White | 0.60 [0.49] | | 53,232 | 0.61 [0.49] | | 53,232 | -0.014 (0.0030) | -2.34 |
| Hispanic/Latino | 0.16 [0.37] | | 53,232 | 0.13 [0.34] | | 53,232 | 0.029 (0.0022) | 18.0 |
| Disability Status | 0.029 [0.17] | | 53,232 | 0.025 [0.16] | | 53,232 | 0.0036 (0.00099) | 12.7 |
| No HS Diploma/GED | 0.034 [0.18] | | 53,232 | 0.028 [0.16] | | 53,232 | 0.0056 (0.0011) | 16.6 |
| HS Diploma/GED | 0.42 [0.49] | | 53,232 | 0.43 [0.49] | | 53,232 | -0.012 (0.0030) | -2.84 |
| Some College | 0.20 [0.40] | | 53,232 | 0.18 [0.39] | | 53,232 | 0.011 (0.0024) | 5.77 |
| College Degree Plus | 0.35 [0.48] | | 53,232 | 0.35 [0.48] | | 53,232 | -0.0031 (0.0029) | -0.89 |
| B. Social Benefits | | | | | | | | |
| Disadvantaged Adult Status | 0.48 [0.50] | | 53,232 | 0.42 [0.49] | | 53,232 | 0.062 (0.0030) | 12.9 |
| Dislocated Worker Status | 0.52 [0.50] | | 53,232 | 0.52 [0.50] | | 53,232 | 0.0055 (0.0031) | 1.06 |
| Low-income Status | 0.42 [0.49] | | 53,232 | 0.40 [0.49] | | 53,232 | 0.024 (0.0030) | 5.75 |
| TANF Recipient | 0.0094 [0.096] | | 53,232 | 0.0077 [0.087] | | 53,232 | 0.0017 (0.00056) | 17.7 |
| SNAP Recipient | 0.21 [0.41] | | 37,768 | 0.18 [0.38] | | 38,114 | 0.031 (0.0029) | 14.8 |
| Other Public Assistance Recipient | 0.25 [0.43] | | 41,854 | 0.24 [0.43] | | 38,234 | 0.0084 (0.0030) | 3.35 |
| C. Pre-Separation Employment | | | | | | | | |
| Employed at Participation | 0.30 [0.46] | | 53,232 | 0.20 [0.40] | | 53,232 | 0.098 (0.0026) | 33.2 |
| Real Earnings 1Q Before | 10687.2 [11522.8] | | 53,232 | 10813.2 [10387.7] | | 53,232 | -126.0 (67.2) | -1.18 |
| Real Earnings 2Q Before | 11139.5 [9657.6] | | 53,232 | 11311.5 [10290.1] | | 53,232 | -172.0 (61.2) | -1.54 |
| Real Earnings 3Q Before | 11019.9 [9689.5] | | 53,232 | 11015.4 [9222.9] | | 53,232 | 4.52 (58.0) | 0.041 |
| Quarters in WIOA/WP | 5.05 [3.43] | | 53,232 | 3.15 [2.58] | | 53,232 | 1.90 (0.019) | 37.6 |

Notes: Columns (1) and (2) show summary statistics for the job trainees with high AI exposure prior to participation while columns (3) and (4) present the same for nearest neighbor matched Wagner-Peyser participants. T-test standard errors in column (5) are adjusted for heteroskedasticity.

Table A.5: Balance Table for Matched Pairs with Low AI Exposure Prior to Participation

| | Trainees | | | Matched Wagner Peyser | | | Trainees - Wagner Peyser | |
|-------------------------------------|---------------------|---------------------|-----------------|-----------------------|---------------------|-----------------|--------------------------|--------|
| | Mean/SD (1) | Δ /SE (2) | # Spells (3) | Mean/SD (4) | Δ /SE (5) | # Spells (6) | Δ /SE | % Diff |
| A. Demographic | | | | | | | | |
| Female | 0.44 [0.50] | | 54,983 | 0.46 [0.50] | | 54,983 | -0.015 (0.0030) | -3.35 |
| Age | 37.8 [10.7] | | 54,983 | 39.2 [10.8] | | 54,983 | -1.39 (0.065) | -3.68 |
| Asian | 0.027 [0.16] | | 54,983 | 0.023 [0.15] | | 54,983 | 0.0046 (0.00094) | 16.9 |
| Black | 0.29 [0.45] | | 54,983 | 0.31 [0.46] | | 54,983 | -0.022 (0.0028) | -7.71 |
| White | 0.57 [0.49] | | 54,983 | 0.57 [0.49] | | 54,983 | 0.00091 (0.0030) | 0.16 |
| Hispanic/Latino | 0.15 [0.36] | | 54,983 | 0.12 [0.32] | | 54,983 | 0.035 (0.0020) | 22.9 |
| Disability Status | 0.025 [0.16] | | 54,983 | 0.021 [0.14] | | 54,983 | 0.0038 (0.00090) | 15.4 |
| No HS Diploma/GED | 0.055 [0.23] | | 54,983 | 0.046 [0.21] | | 54,983 | 0.0092 (0.0013) | 16.7 |
| HS Diploma/GED | 0.52 [0.50] | | 54,983 | 0.56 [0.50] | | 54,983 | -0.037 (0.0030) | -7.19 |
| Some College | 0.19 [0.39] | | 54,983 | 0.18 [0.38] | | 54,983 | 0.015 (0.0023) | 7.65 |
| College Degree Plus | 0.23 [0.42] | | 54,983 | 0.22 [0.41] | | 54,983 | 0.015 (0.0025) | 6.53 |
| B. Social Benefits | | | | | | | | |
| Disadvantaged Adult Status | 0.58 [0.49] | | 54,983 | 0.50 [0.50] | | 54,983 | 0.081 (0.0030) | 14.0 |
| Dislocated Worker Status | 0.42 [0.49] | | 54,983 | 0.39 [0.49] | | 54,983 | 0.026 (0.0030) | 6.23 |
| Low-income Status | 0.46 [0.50] | | 54,983 | 0.43 [0.50] | | 54,983 | 0.031 (0.0030) | 6.65 |
| TANF Recipient | 0.0092 [0.095] | | 54,983 | 0.0072 [0.084] | | 54,983 | 0.0020 (0.00054) | 21.8 |
| SNAP Recipient | 0.23 [0.42] | | 40,865 | 0.18 [0.38] | | 41,285 | 0.052 (0.0028) | 22.6 |
| Other Public Assistance Recipient | 0.27 [0.44] | | 45,662 | 0.26 [0.44] | | 40,273 | 0.013 (0.0030) | 4.85 |
| C. Pre-Separation Employment | | | | | | | | |
| Employed at Participation | 0.37 [0.48] | | 54,983 | 0.28 [0.45] | | 54,983 | 0.084 (0.0028) | 22.9 |
| Real Earnings 1Q Before | 9287.2 [10355.7] | | 54,983 | 9418.4 [10874.9] | | 54,983 | -131.2 (64.0) | -1.41 |
| Real Earnings 2Q Before | 9966.1 [9522.3] | | 54,983 | 10005.0 [10447.8] | | 54,983 | -38.9 (60.3) | -0.39 |
| Real Earnings 3Q Before | 9785.7 [9028.0] | | 54,983 | 9715.5 [10259.7] | | 54,983 | 70.2 (58.3) | 0.72 |
| Quarters in WIOA/WP | 4.60 [3.23] | | 54,983 | 2.94 [2.51] | | 54,983 | 1.66 (0.017) | 36.1 |

Notes: Columns (1) and (2) show summary statistics for the job trainees with high AI exposure prior to participation while columns (3) and (4) present the same for nearest neighbor matched Wagner-Peyser participants. T-test standard errors in column (5) are adjusted for heteroskedasticity.

Table A.6: Regression Estimates for Quarterly Earnings Returns to WIOA/WIA Training

| | Prior AI Exposure | | Target Occ. AI Exposure | |
|---------------------------------|-------------------|------------------|-------------------------|------------------|
| | High | Low | High | Low |
| A. Pooled Post: | 1,466 (51.96) | 1,945 (47.94) | 1,039 (49.52) | 2,272 (49.65) |
| B. Pooled Pre (Placebo): | -80 (44.32) | -7 (48.63) | 91 (51.24) | -99 (41.33) |
| N | 53,232 | 54,983 | 47,358 | 60,700 |
| Mean Trainee Baseline Wages: | 10,949 | 9,680 | 11,530 | 9,340 |

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.7: Headline Results for Share of Participants and Occupations “AI-Retrainable”

| AI Retraining Statistic | | Trainees | | Trainees - Matched | |
|-------------------------|--|-------------|--------------------------|--------------------|--------------------------|
| | | Full Sample | High Initial AI Exposure | Full Sample | High Initial AI Exposure |
| | | (1) | (2) | (3) | (4) |
| Sh. Spells | $\Delta SML \geq 0 \ \& \ \Delta Y \geq 0$ | 35.9% | 21.1% | 28.8% | 25.9% |
| Sh. Spells | $\Delta SML^{LLM} \geq 0 \ \& \ \Delta Y \geq 0$ | 41.2% | 23.6% | 34.4% | 27.7% |
| Sh. Occups | $\Delta \overline{SML} \geq 0 \ \& \ \Delta \overline{Y} \geq 0$ | 35.7% | 11.4% | 26.2% | 20.8% |
| Sh. Occups | $\Delta \overline{SML}^{LLM} \geq 0 \ \& \ \Delta \overline{Y} \geq 0$ | 38.4% | 8.8% | 39.1% | 22.8% |

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. SML corresponds to the Brynjolfsson et al. (2018) AI exposure measure while SML^{LLM} corresponds to the Eloundou et al. (2024) measure. A 6-digit occupation is labeled AI-retrainable if the mean worker in that occupations AI-retrainable. See text for further details.

Appendix B. 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

$$Y_i = D_i Y_i(1, H_i) + (1 - D_i) Y_i(0, H_i). \quad (\text{A.1})$$

We want group-specific causal ATE parameters:

$$\tau_h \equiv \mathbb{E}[Y_i(1, h) - Y_i(0, h) \mid H_i = h], \quad h \in \{0, 1\}. \quad (\text{A.2})$$

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

$$\delta_h \equiv \mathbb{E}[Y_i(1, h) \mid D_i = 1] - \mathbb{E}[Y_i(0, h) \mid D_i = 0]. \quad (\text{A.3})$$

Lingering Selection Bias within Each Group

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

$$\delta_h = \underbrace{\tau_h}_{\text{causal ATE}} + \underbrace{\left(\mathbb{E}[Y_i(0, h) \mid D_i = 1] - \mathbb{E}[Y_i(0, h) \mid D_i = 0] \right)}_{B_h \text{ (lingering selection bias within } h)}. \quad (\text{A.4})$$

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:

$$\Delta \equiv \delta_1 - \delta_0 = \underbrace{(\tau_1 - \tau_0)}_{\text{heterogeneous effect of interest}} + \underbrace{(B_1 - B_0)}_{\text{residual/differential bias}}. \quad (\text{A.5})$$

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

$$B_h = \bar{B} + \varepsilon_h, \quad (\text{A.6})$$

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

$$\Delta - (\tau_1 - \tau_0) = (B_1 - B_0) = \varepsilon_1 - \varepsilon_0, \quad (\text{A.7})$$

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).

$$\mathbb{E}[Y_i(0, h) \mid D_i = 1, H_i = h] - \mathbb{E}[Y_i(0, h) \mid D_i = 0, H_i = h] = \bar{B} \quad \text{for } h \in \{0, 1\}, \quad (\text{A.8})$$

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 C. Matching Estimator Details

We use the semiparametric 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. We consider a distance metric $\|X_i - X_j\|$ that creates a score between each trainee's covariates X_i and all potential j control candidates with covariate vectors X_j in the years prior to WIOA/WIA training, and keep the closest match for each treated unit i . To equalize scales among the different components of each vector in their contribution to the distance score, we use a Mahalanobis distance statistic that weights vector distances by the inverse of the variance-covariance matrix (Σ) for covariates:

$$\|X_i - X_j\|_{Mahalanobis} = \sqrt{(X_i - X_j)'(\Sigma_x^{-1})(X_i - X_j)} \quad (\text{A.9})$$

We use the algorithm discussed in [Abadie et al. \(2004\)](#) to implement this procedure as it has the useful property that violators of the “covariate overlap” assumption are dropped, resulting in treated units only being used if the covariates share a sufficiently common covariate support. In practice, less than 1% of our final sample is dropped due to overlap, a high match success rate that arises due to the abundance of potential control units available in the Wagner-Peyser program.