

The Labour Market Effects of AI-Related Hiring Intensity

by

Fatemeh Kalantari

An Extended Essay Submitted in Partial Fulfillment of the

Requirements for the Degree of

MASTER OF ARTS

in the Department of Economics



Dr. Tao Wang, Supervisor (Department of Economics)

Dr. Kenneth G. Stewart, Member (Department of Economics)

© Fatemeh Kalantari, 2026

University of Victoria

All rights reserved. This essay may not be reproduced in whole or in part,

by photocopying or other means, without the permission of the author.

Abstract

This paper investigates how national AI-related hiring intensity affects wages and employment across US industries with different levels of AI exposure. Using data on 145 four-digit NAICS industries from 2019 to 2025, I estimate a continuous Difference-in-Differences model that interacts each industry's AI exposure level — measured by the AI Industry Exposure (AIIE) index from Felten, Raj, and Seamans (2021), which captures structural task exposure to AI capabilities rather than actual AI adoption — with the national share of US Indeed job postings requiring AI skills as a time-varying measure of AI-related hiring intensity.

I find three main results. First, rising AI-related hiring intensity has no significant effect on average weekly wages in any industry group, suggesting that wage adjustment to AI technology is slow and has not yet materialised in the data. Second, within the low-exposure industry group, a one unit increase in the full-sample standardised AIIE index is associated with a 3.8 percent differential employment gain for a one percentage point increase in national AI-related hiring intensity ($\beta = +0.038$, $p < 0.05$), consistent with a labour reallocation hypothesis in which employment growth is stronger in industries less directly substitutable by AI technology. Third, a supplementary analysis using the ChatGPT launch as a discrete treatment event (Post = 1 from Q1 2023, the first observed post-launch quarter in the balanced employment panel) finds a significant employment gain within the low-exposure tercile ($\beta = +0.058$, $p < 0.05$), suggesting reallocation dynamics intensified in the post-generative-AI period. The results suggest that rising national AI-related hiring intensity has not yet produced measurable wage divergence across industries, but is associated with employment reallocation patterns within the low-exposure segment of the industry distribution. These findings are consistent with recent evidence showing null aggregate wage effects of generative AI (Kauhanen and Rouvinen, 2025; Hartley et al., 2024; Massenkoff and McCrory, 2026).

Keywords: AI-related hiring intensity, Artificial intelligence, Difference-in-Differences, employment, industry exposure, labour reallocation, wages

Contents

Abstract	ii
List of Tables	v
List of Figures	vi
1 Introduction	1
2 Literature Review	7
2.1 Technology, Tasks, and Labour Market Polarisation	7
2.2 Automation, Robots, and Displacement.....	8
2.3 Measuring AI Exposure.....	9
2.4 Empirical Evidence on AI and Labour Market Outcomes	11
3 Data	14
3.1 Wage Data — QCEW Monthly Panel	14
3.2 Employment Data — QCEW Quarterly Panel.....	14
3.3 AI Exposure Index — AIIE (Felten et al., 2021)	15
3.4 AI-Related Hiring Intensity — Indeed Hiring Lab	15
3.5 Data Construction and Summary Statistics	16
4 Empirical Framework	19
4.1 Core Continuous DiD Models	19
4.2 Identification and Fixed Effects	20

4.3	Parallel Trends Testing.....	21
4.4	Methodological Notes.....	21
5	Results	23
5.1	Parallel Trends — Event Studies	23
5.2	Core Results: Wages	26
5.3	Core Results: Employment	28
5.4	Supplementary: ChatGPT Binary DiD.....	38
6	Discussion	41
7	Conclusion	43
A	COVID Extension	48

List of Tables

1	Data Construction — Sample Selection.....	17
2	Summary Statistics — Main Variables.....	18
3	Core DiD — Effect of AI-Related Hiring Intensity on Log Weekly Wage.....	26
4	Core DiD — Effect of AI-Related Hiring Intensity on Log Employment.....	29
5	Supplementary — ChatGPT Binary DiD — Wages and Employment	39
A1	COVID Extension — Effect of AI Exposure on Wages and Employment	49
A2	COVID Extension with Remote Work Control (Dingel and Neiman, 2020)	50

List of Figures

1	Event Study — Wage (Monthly)	23
2	Event Study — Wage (Binned Periods).....	24
3	Event Study — Employment (Quarterly)	25
4	Average Log Weekly Wage by AI Exposure Group	28
5	Average Log Employment by AI Exposure Group	30
6	Employment Change by Industry — Raw Top Gainers and Losers	32
7	Employment Change Relative to Industry Fixed Effects — Post vs Pre-ChatGPT	34
8	Employment Trends by Exposure Group and Within-Group AIIIE	35
9	Employment Change vs Exposure — Low Exposure Only	36
10	Employment Change vs AI Exposure by Tercile	37
A1	COVID Heterogeneity by Exposure Tercile	51

1 Introduction

Artificial intelligence is increasingly recognised as one of the most transformative forces reshaping labour markets in the modern era. The task-based framework established by Autor et al. (2003) showed that technological change substitutes for routine tasks while complementing non-routine cognitive and manual work — a framework that has guided empirical research on automation for two decades. Building on this, Acemoglu and Restrepo (2018) developed a theoretical model in which automation creates displacement effects in directly affected industries, and Acemoglu and Restrepo (2020) provided empirical support for this mechanism, finding that industrial robots reduced both employment and wages in exposed US commuting zones. These foundational contributions established that technology’s labour market effects are heterogeneous — they depend critically on the task content of affected work rather than on broad industry classifications.

What distinguishes artificial intelligence from earlier waves of automation is the nature of the tasks it can perform. Frey and Osborne (2017) estimated that approximately 47 percent of US employment is at high risk of computerisation, while Brynjolfsson et al. (2018) showed using a task-level rubric applied to the O*NET database that machine learning affects a fundamentally different set of occupations than earlier waves of automation, with few occupations fully automatable and most requiring a redesign of job tasks rather than outright replacement of workers. Unlike industrial robots that primarily displaced manufacturing workers, AI threatens occupations that previous automation left largely untouched — lawyers, financial analysts, writers, and software developers. This raises the possibility that AI may produce a fundamentally different pattern of labour market disruption than prior technologies, one that cuts across the skill distribution rather than concentrating in routine middle-skill occupations as earlier automation did.

The theoretical literature offers two competing channels through which this disruption could

operate. Acemoglu and Restrepo (2019) formalise a displacement effect, whereby automation substitutes directly for labour in affected tasks and reduces demand for the workers who previously performed them, alongside a reinstatement effect, whereby the productivity gains from automation create entirely new tasks and occupations that require human labour in complementary roles. The net effect of AI on any given industry's employment and wages depends on the relative strength of these two forces, and theory alone does not predict which will dominate. This ambiguity is precisely what motivates an empirical industry-level investigation: if displacement dominates in the most AI-exposed industries, employment should contract there relative to less-exposed industries; if reinstatement dominates, the opposite pattern should emerge, or AI exposure may correlate with stronger employment growth as exposed industries expand into newly created tasks.

This task-based perspective draws an important distinction between manual and cognitive work that is central to understanding AI's likely labour market footprint. Autor et al. (2003) classify occupational tasks along two dimensions: routine versus non-routine, and manual versus cognitive. Industrial-era automation, including the industrial robots studied by Acemoglu and Restrepo (2020), primarily substituted for routine manual tasks — assembly, machine operation, and other physically repetitive work concentrated in manufacturing. Acemoglu and Autor (2011) review evidence consistent with this substitution pattern as a driver of labour market polarisation, as routine manual and routine cognitive occupations alike were hollowed out while non-routine work at both ends of the skill distribution expanded. Artificial intelligence breaks from this pattern by exhibiting little capacity to substitute for non-routine manual tasks — physical labour requiring dexterity, mobility, or hands-on judgement, such as construction, repair, and many forms of healthcare and personal services — while exhibiting substantial capacity to substitute for non-routine cognitive tasks, including the writing, analysis, and communication-based work performed by many professional and white-collar occupations. Webb (2020) confirms this empirically, showing that AI exposure is concentrated among high-skill, high-wage occupations rather than the routine manual occupations most affected by earlier automation. This asymmetry motivates a central distinction in this paper's empirical design: industries whose work is structurally manual and non-routine are classified

as low-exposure, while industries whose work is structurally cognitive and language-based are classified as high-exposure, following the exposure measure of Felten et al. (2021).

The rapid diffusion of generative AI tools following the launch of ChatGPT in November 2022 has given new urgency to these theoretical questions. Eloundou et al. (2023) find that large language models have the potential to affect a substantial share of all US occupational tasks, with higher-income occupations facing greater exposure. Hui et al. (2024) provide early empirical evidence from an online labour market, finding that freelancers in directly affected occupations experienced significant reductions in both employment and earnings following ChatGPT’s release. Brynjolfsson et al. (2023) show in a field experiment that generative AI tools significantly increased productivity among customer service workers, with the largest gains among lower-skill employees. Yet this emerging evidence base remains fragmented along two dimensions. First, much of it is drawn from micro-level settings — individual firms, freelance platforms, or controlled experiments — that, while informative about specific mechanisms, are not designed to speak to economy-wide patterns of wage and employment adjustment. Second, where macro-level or national evidence exists, it has tended to focus on either wages or employment in isolation, rather than examining both outcomes jointly within a single empirical framework. Despite this growing evidence, relatively little is known about how AI-related hiring intensity affects wages and employment together at the industry level across the broader US economy.

A small but growing body of recent work has begun to examine these aggregate questions directly, with results that are so far decidedly mixed. Kauhanen and Rouvinen (2025) find no statistically significant changes in wages or employment in Finnish administrative data through 2024, while Hartley et al. (2024) report similarly small wage effects using US occupational data. Humlum and Vestergaard (2025) use matched employer-employee survey data from Denmark to rule out earnings and hours effects larger than two percent two years after ChatGPT’s launch, even among intensive AI users. Against this, Brynjolfsson et al. (2025) find using high-frequency US payroll data that early-career workers in AI-exposed occupations have experienced relative employment

declines of approximately sixteen percent since late 2022, even as employment for more experienced workers in the same occupations has remained stable. This divergence between near-zero aggregate effects and substantial effects concentrated among specific worker groups suggests that the labour market consequences of AI may be obscured when studied only in aggregate, or only among one segment of the workforce — motivating the industry-level, exposure-based approach taken in this paper.

An industry-level approach offers a useful middle ground between micro-level case studies and purely macroeconomic aggregates. Industries aggregate many individual firms and workers, smoothing out idiosyncratic shocks while retaining enough granularity — in this paper, 145 four-digit NAICS industries — to exploit meaningful variation in the structural task content of work. This level of analysis also aligns naturally with the publicly available data infrastructure maintained by US statistical agencies, allowing the results to be fully reproducible and extended by future researchers as additional post-ChatGPT data become available. Beyond its methodological convenience, the industry level is also where much of the current policy debate is situated: discussions of retraining programmes, sectoral labour market policy, and industrial strategy in response to AI typically reference industries or broad occupational groups rather than individual firms or workers, making industry-level evidence directly relevant to the design of labour market policy.

This paper contributes to filling that gap. I estimate the effect of national AI-related hiring intensity on wages and employment across US industries using a continuous Difference-in-Differences design. The empirical strategy interacts an industry-level AI exposure index — the AIIE index of Felten et al. (2021) — with the national share of US job postings requiring AI skills from Indeed’s Hiring Lab. Since $AIshare_t$ varies only over time and is identical across all industries in any given period, and since $Exposure_i$ varies only across industries and is fixed over time, the interaction term varies because the same national time shock is loaded differently by industries with different levels of AI exposure. Industries with a higher AIIE score amplify the national trend more than industries with a lower score. This follows an exposure-based shift-share logic — industries

with different pre-determined AI exposure are differentially loaded onto the same national AI-hiring shock — following the framework of Goldsmith-Pinkham et al. (2020).

Johnston and Makridis (2025) provide the most directly comparable evidence to date, finding that a one standard deviation increase in industry-level AI exposure raises output by approximately seven percent, with positive employment effects concentrated in industries where AI is more likely to complement rather than substitute for human labour. Their results suggest that the net employment effect of AI exposure depends on the underlying task structure of the industries involved, rather than being uniformly positive or negative — a finding consistent with the displacement and reinstatement channels discussed above. The paper extends their framework in three important ways. First, it incorporates employment as an additional outcome alongside wages, directly addressing the gap noted above between studies that examine only one margin of adjustment. Second, it uses a longer and richer panel covering 2019 to 2025, spanning the pre-pandemic period, the COVID-19 shock, and the full post-ChatGPT diffusion of generative AI. Third, it treats the ChatGPT launch as a supplementary discrete treatment event, allowing a comparison between the continuous national hiring-intensity measure and a sharper, more widely recognised structural break in AI capabilities. A COVID extension is included in Appendix A both to provide a fuller picture of the sample period and to demonstrate that the COVID-19 shock, while economically important, does not drive the paper’s central findings.

The results show that rising AI-related hiring intensity has no significant differential effect on wages. However the employment analysis reveals important heterogeneity: within the low-exposure industry group, those with relatively higher AIIE scores are associated with significant employment gains as national AI-related hiring intensity rises ($\beta = +0.038$, $p < 0.05$). A supplementary ChatGPT binary DiD confirms this pattern ($\beta = +0.058$, $p < 0.05$). Together these results are consistent with a labour reallocation interpretation in which the reinstatement channel described above operates most visibly not within the most AI-exposed industries themselves, but in the relatively less-exposed industries that appear to absorb part of the resulting shift in labour demand.

The remainder of the paper proceeds as follows. Section 2 reviews the literature. Section 3 describes the data. Section 4 presents the empirical framework. Section 5 reports the results. Section 6 discusses the findings. Section 7 concludes. The COVID extension is in Appendix A.

2 Literature Review

2.1 Technology, Tasks, and Labour Market Polarisation

The modern literature on technology and labour markets is built on the task-based framework of Autor et al. (2003). They show that computers substitute for routine cognitive and manual tasks while complementing non-routine tasks requiring judgment and problem-solving. This framework was extended directly to UK data by Goos and Manning (2007), who document that the United Kingdom has exhibited a pattern of job polarisation since 1975, with employment shares rising at both the highest- and lowest-wage occupations while middle-wage occupations contracted. They argue that this pattern is better explained by the routinisation hypothesis of Autor et al. (2003) than by a simple skill-biased technical change story, and show that job polarisation alone accounts for roughly one-third of the rise in the gap between median and low-wage earnings and about one-half of the rise in the gap between high and median earnings over their sample period. Acemoglu and Autor (2011) generalise this task-based framework theoretically, formalising how technology, trade, and offshoring jointly determine the assignment of tasks to labour and capital across the skill distribution, providing the theoretical backbone for interpreting polarisation as a structural rather than purely cyclical phenomenon.

Autor (2019) extends this line of inquiry geographically, showing that US cities have grown markedly more skill-intensive over the past five decades, yet non-college workers within those cities have been pushed out of specialised middle-skill occupations into lower-wage jobs requiring only generic skills — a process he attributes jointly to automation and trade, and which has eroded the wage premium that non-college workers in dense urban labour markets once commanded. Deming (2017) complements this evidence by showing that social skills have become increasingly valuable in the labour market as technology advances: between 1980 and 2012, jobs requiring high levels

of social interaction grew substantially while routine-intensive jobs shrank. Deming’s explanation centres on a model of team production in which workers specialise according to comparative advantage; because interpersonal coordination, persuasion, and judgment remain difficult for machines to replicate, workers who supply these skills become more valuable precisely as routine tasks are increasingly automated. Taken together, this literature establishes that technological change reshapes labour markets unevenly across the skill and task distribution, not uniformly across industries or occupations — a premise that motivates the structural, task-based exposure measure used later in this paper.

2.2 Automation, Robots, and Displacement

Acemoglu and Restrepo (2018) formalise these dynamics in an explicit theoretical model, distinguishing a displacement effect, in which automation substitutes directly for labour in the tasks it newly performs, from a productivity effect, in which the resulting cost savings raise output and can increase the demand for labour in tasks that remain non-automated. Whether automation depresses or raises aggregate employment and wages in their model depends on the relative strength of these two forces, and on whether new tasks are created for labour to perform as old ones are automated — a tension that anticipates the displacement-reinstatement framework discussed further below. Acemoglu and Restrepo (2020) test this model empirically using the adoption of industrial robots across US commuting zones, finding that robot exposure reduced both employment and wages, with each additional robot per thousand workers associated with an employment decline of approximately 0.2 percentage points; their results imply that the displacement effect dominated the productivity effect for robots over their sample period, at least within the directly affected local labour markets. Graetz and Michaels (2018) examine a similar question using panel data on industrial robot adoption across 17 countries between 1993 and 2007, finding a markedly different aggregate picture: increased robot use contributed approximately 0.36 percentage points to annual labour productivity growth and raised both total factor productivity and wages, while lowering output prices, and did not significantly reduce total employment — although it did reduce the employment

share of low-skilled workers specifically. The contrast between the local, commuting-zone-level displacement documented by Acemoglu and Restrepo (2020) and the more favourable aggregate, cross-country productivity and wage effects in Graetz and Michaels (2018) illustrates a recurring theme in this literature: automation’s costs are often concentrated and highly visible locally, among directly displaced workers and their communities, even when its aggregate productivity benefits are diffuse and broadly positive.

Webb (2020) extends this body of work to artificial intelligence specifically, constructing exposure measures for robots, software, and AI based on the overlap between patent text and detailed occupational task descriptions. He first validates the method historically, showing that occupations highly exposed to software and industrial robots experienced employment and wage declines over the relevant periods, before applying the same approach to AI. In sharp contrast to software and robots, he finds that AI exposure is directed at high-skilled tasks: the most AI-exposed occupations include clinical laboratory technicians, chemical engineers, optometrists, and power plant operators, and older, more experienced workers are substantially more exposed to AI than younger workers — the opposite of the pattern for robots, whose exposure is concentrated among workers without a high school diploma and declines monotonically with education. This asymmetry — AI threatening occupations and worker groups that prior waves of automation left comparatively untouched — is central to the motivation for treating AI as a distinct technological shock in the present paper rather than as a continuation of the robotics literature.

2.3 Measuring AI Exposure

A central methodological challenge in this literature is translating the conceptual distinction between routine and non-routine, or manual and cognitive, tasks into a measurable index of technological exposure at the occupation or industry level. Frey and Osborne (2017) took an early and influential approach, asking a panel of machine learning researchers to assess the automatability of a sample of occupations based on engineering bottlenecks — tasks requiring fine motor control, creativity, or social intelligence were judged harder to automate — and using these assessments to

train a classifier applied to the full set of US occupations in the O*NET database. Their resulting estimate, that approximately 47 percent of US employment is at high risk of computerisation, became one of the most widely cited figures in the automation debate, though subsequent work has often found this figure to overstate near-term displacement risk by treating occupations rather than tasks as the unit of automation. Brynjolfsson et al. (2018) take a more granular, task-level approach, applying a rubric for "Suitability for Machine Learning" to thousands of individual work activities drawn from O*NET. They find that machine learning is likely to affect a fundamentally different set of occupations than earlier waves of automation, that few occupations are fully automatable even where some constituent tasks are highly suitable for machine learning, and that realising machine learning's productivity potential typically requires redesigning job tasks rather than simply replacing workers wholesale.

Felten et al. (2021), whose index is used in this paper, build on this task-level tradition but focus specifically on artificial intelligence rather than machine learning broadly. Their AI Industry Exposure (AIIE) index is constructed by first surveying respondents, via Amazon Mechanical Turk, on how relevant each of ten distinct AI application areas is to each of fifty-two human abilities catalogued in O*NET, producing an application-by-ability relatedness matrix; this matrix is then combined with each occupation's own O*NET ability profile and aggregated to the four-digit NAICS industry level using BLS employment weights. The resulting AIIE score is the most comprehensive industry-level AI exposure measure currently available, and importantly captures structural exposure to AI capabilities based on task content — it does not indicate that an industry has actually adopted AI more intensively than others, a distinction that is central to the identification strategy developed in Section 4. In their own data, Felten et al. (2021) report that AI exposure is highest among service industries characterised by intensive information processing, while the lowest-scoring industries are blue-collar sectors involving manual labour, such as crop production support services, construction contracting, and warehousing and storage — precisely the kind of industries that this paper's own results, discussed in Section 5.3, identify as employment gainers within the low-exposure tercile. Eloundou et al. (2023) provide a complementary measure specific

to large language models, estimating that GPT-class models can affect at least 10 percent of tasks in approximately 80 percent of US occupations, and that this exposure rises further when GPT-powered software tools, rather than the models themselves, are taken into account. Relative to the broad computerisation-risk measure of Frey and Osborne (2017) and the machine-learning-specific rubric of Brynjolfsson et al. (2018), the AIIE index’s grounding in ten distinct AI application areas and its industry-level aggregation make it particularly well suited to the continuous, industry-level empirical design adopted in this paper.

2.4 Empirical Evidence on AI and Labour Market Outcomes

Several studies find null or small aggregate effects. Kauhanen and Rouvinen (2025) use Finnish administrative records finding no statistically significant changes in wages or employment through 2024. Hartley et al. (2024) similarly find only small wage effects in US occupational data. Massenkoff and McCrory (2026) use actual Claude usage data from Anthropic, finding no systematic unemployment increase since late 2022. Other studies find more significant effects at the micro level. Hui et al. (2024) find significant employment and earnings reductions for freelancers following ChatGPT. Brynjolfsson et al. (2023) find that generative AI raised productivity most among lower-skill workers. Noy and Zhang (2023) find in a controlled experiment that generative AI tools significantly raised productivity in writing tasks, with the largest gains among lower-skill workers. Johnston and Makridis (2025), the closest antecedent to this paper, find using proprietary administrative data that a one standard deviation increase in industry AI exposure raises output by 7 percent and employment by 4 percent where AI requires human collaboration. The present paper builds on their approach using publicly available QCEW data with a time-varying national treatment measure and four-digit NAICS granularity.

More recent evidence strengthens the null aggregate findings while identifying concentrated effects among specific worker groups. Humlum and Vestergaard (2025) use two large-scale adoption surveys linked to matched employer-employee records in Denmark across 11 AI-exposed occupations, estimating precise null effects on earnings and hours at both the worker and workplace level

— ruling out effects larger than two percent two years after ChatGPT’s launch. These null results hold even for intensive AI users, early adopters, and workplaces with substantial AI investments, providing some of the strongest evidence yet against near-term wage disruption from generative AI. In contrast, Brynjolfsson et al. (2025) use high-frequency administrative payroll data from ADP covering millions of US workers and find that early-career workers aged 22 to 25 in AI-exposed occupations experienced approximately 16 percent relative employment declines since late 2022, while employment for experienced workers remained stable. These employment adjustments occur via headcount rather than compensation and are concentrated in occupations where AI automates rather than augments labour. This pattern — employment adjusting before wages — is directly consistent with the null wage and significant employment reallocation findings reported in the present paper. PwC’s (2025) Global AI Jobs Barometer, based on close to one billion job advertisements across six continents, finds that wages are growing twice as fast in industries more exposed to AI and that AI-skilled workers command a 56 percent wage premium over otherwise similar workers. This appears to contrast with the null wage result here, but the difference is interpretable — PwC measures the wage premium for possessing AI skills within occupations, while this paper estimates industry-level wage divergence conditional on fixed effects — these are measuring distinct phenomena.

The theoretical mechanism connecting AI adoption to employment reallocation is developed most rigorously by Acemoglu and Restrepo (2019), who distinguish between the displacement effect — automation substitutes for labour in directly affected tasks — and the reinstatement effect — new tasks emerge that require human labour in complementary roles. When automation raises productivity in high-exposure sectors, the resulting output gains generate demand for goods and services produced by less-exposed industries, shifting labour demand toward those sectors. This reallocation mechanism predicts that employment losses in high-exposure industries should be offset by gains elsewhere in the economy — not necessarily in the industries most complementary to AI, but in industries that absorb the reallocated labour force. The present paper tests whether this reallocation is visible in industry-level employment data by examining whether industries with

relatively higher AI exposure within the low-exposure group capture disproportionate employment gains as national AI-related hiring intensity rises.

3 Data

3.1 Wage Data — QCEW Monthly Panel

Wage data are drawn from the Bureau of Labor Statistics Quarterly Census of Employment and Wages (QCEW), which covers approximately 95 percent of all US wage and salary workers. Although the QCEW publishes aggregate wage totals at the quarterly level, the monthly average weekly wage figures used here are sourced from the QCEW monthly microdata files, which record payroll submissions separately for each reference month within each quarter.

The raw wage data covers January 2019 to December 2025 across 145 four-digit NAICS industries — a span of 84 possible months. December 2025 observations are excluded from the final panel because the corresponding AI posting share from Indeed Hiring Lab was not yet published at the time of data collection, reducing the effective sample to 83 months running from January 2019 to November 2025. After merging with the Indeed AI posting share and removing suppressed observations, the final wage panel contains 11,295 observations. The panel is unbalanced at the monthly level as some smaller industries have suppressed wage observations in certain months in the underlying QCEW microdata, resulting in effective coverage of approximately 78 months per industry out of the maximum of 83. All wage models use the natural logarithm of average weekly wages as the outcome variable. Because the CPI price index varies only over time it is fully absorbed by the monthly time fixed effects, meaning the DiD coefficients are equivalently interpretable as real wage effects.

3.2 Employment Data — QCEW Quarterly Panel

Employment data are also sourced from the QCEW, using quarterly employment counts at the four-digit NAICS industry level measured as the average of monthly employment levels across

the three months of each quarter. The employment sample spans Q1 2019 to Q3 2025 — the panel was extended back to 2019 by downloading additional QCEW quarterly files, providing a proper pre-COVID baseline for the COVID extension in Appendix A. Although the sample spans Q1 2019 to Q3 2025, two quarters — Q4 2021 and Q4 2022 — are excluded because insufficient industries reported employment data in those periods in the downloaded QCEW files. As a consequence, the supplementary ChatGPT binary DiD model defines the post-treatment period as beginning in Q1 2023 rather than Q4 2022. The final balanced panel therefore covers 25 of the 27 possible quarters in the sample window. After merging with the wage data and restricting to industries present in all 25 quarters, the employment panel contains 3,125 observations across 125 industries.

3.3 AI Exposure Index — AIIE (Felten et al., 2021)

The industry-level AI exposure measure is the AIIE index from Felten et al. (2021). The index links progress across ten AI application areas to occupational ability requirements from O*NET, aggregated to the four-digit NAICS level using BLS employment weights. A higher AIIE score reflects industries whose workers perform tasks compatible with current AI capabilities. A high AIIE score indicates structural exposure — not actual AI adoption. The index is standardised to mean zero and standard deviation one and is time-invariant, making it not subject to reverse causality concerns.

3.4 AI-Related Hiring Intensity — Indeed Hiring Lab

The time-varying treatment measure is the national share of US job postings requiring AI skills from the Indeed Hiring Lab AI Tracker, published under Creative Commons Attribution 4.0. The series varies only over time and is identical across all industries — no individual industry can meaningfully influence it. This follows an exposure-based shift-share logic — industries with different pre-determined AI exposure are differentially loaded onto the same national AI-hiring shock — following the framework of Goldsmith-Pinkham et al. (2020). The series starts at approximately 1.7 percent in January 2019, rises to approximately 3.2 percent in early 2022,

then declines during the 2022–2023 technology sector contraction before recovering strongly to approximately 3.9 percent by November 2025. The maximum value reported in Table 2 is 3.169 percent, which reflects the quarterly average of the monthly AI share values for Q3 2025 — the average of July, August, and September 2025 — rather than the peak monthly value of 3.881 percent reached in November 2025.

3.5 Data Construction and Summary Statistics

Table 1 presents the step-by-step construction of both panels, showing exactly how the final samples were derived from the raw data. For the wage panel, the main reduction occurs between Step 2 and Step 3 when merging with the Indeed AI posting share — the 125 observations lost correspond to December 2025 wage entries for which the AI posting share was not yet available. For the employment panel, the largest reduction occurs in Step 2 when the raw QCEW data is filtered to private sector four-digit NAICS industries at the US national level, reducing the sample from 9,093 to 3,952 observations across 338 industries. A further reduction in Step 4 imposes the balanced panel restriction, retaining only the 125 industries present in all 25 available quarters and dropping 20 industries that do not appear consistently throughout the sample period.

Table 2 presents descriptive statistics for the main variables used in the analysis. Average weekly wages show substantial variation across industries — ranging from well below the sample mean of \$1,145 to a maximum of \$2,546 — reflecting the wide disparity in pay across four-digit NAICS industries from food preparation to financial services. Employment is highly right-skewed, with a mean of 672,715 workers but a maximum exceeding 11 million, indicating that a small number of very large industries dominate total employment while most industries are considerably smaller. The AIIE index has mean zero and standard deviation one by construction following standardisation, with values ranging from negative to positive reflecting industries below and above the average level of AI task exposure respectively. The AI posting share rises steadily over the sample period, with the interquartile range of 1.765 to 2.478 percent capturing the bulk of the variation used for identification, and a maximum of 3.169 percent corresponding to the Q3 2025

quarterly average — the last quarter in the employment panel.

Table 1: Data Construction — Sample Selection

Panel	Step	Description	Observations	Industries
Wage Panel	1	Raw QCEW monthly	11,420	145
	2	Remove missing wages and AI share	11,420	145
	3	Merge with Indeed AI posting share	11,295	145
FINAL	4	Final wage panel (Jan 2019 – Nov 2025)	11,295	145
Employment Panel	1	Raw QCEW quarterly files 2018–2025	9,093	338
	2	Filter: private sector, 4-digit NAICS, US national	3,952	338
	3	Merge with wage panel on industry × year × quarter	3,375	145
FINAL	4	Balance panel: keep industries in all 25 quarters	3,125	125

Notes: Wage panel: monthly January 2019 to November 2025, 145 industries, 83 months (December 2025 excluded — Indeed AI posting share not yet available at time of data collection). Employment panel: quarterly 2019–2025, 125 industries, 25 quarters (Q4 2021 and Q4 2022 excluded due to insufficient industry coverage). Source: BLS QCEW; Indeed Hiring Lab.

Table 2: Summary Statistics — Main Variables

Variable	N	Mean	SD	P25	P50	P75	Max
Average Weekly Wage (\$)	11,295	1,145	417	828	1,127	1,417	2,546
Log Weekly Wage	11,295	6.972	0.389	6.719	7.027	7.256	7.842
Employment (quarterly avg)	3,125	672,715	1,150,594	145,758	303,101	819,659	11,168,222
AIIE Exposure Index	3,125	0.000	1.000	−0.779	−0.138	0.696	2.218
AI Share Postings (%)	3,125	2.169	0.500	1.765	2.036	2.478	3.169

Notes: Wage panel $N = 11,295$ (145 industries, up to 83 months, January 2019 to November 2025). Employment panel $N = 3,125$ (125 industries, 25 quarters, Q1 2019 to Q3 2025; Q4 2021 and Q4 2022 excluded due to insufficient industry coverage). AIIE = AI Industry Exposure index, standardised to mean = 0, SD = 1 (Felten et al., 2021). P25/P50/P75 = 25th, 50th, and 75th percentiles. AI posting share statistics reflect quarterly averages in the employment panel — the Q3 2025 maximum of 3.169 percent is the average of July, August, and September 2025 monthly values, and is therefore lower than the peak monthly value of 3.881 percent reached in the wage panel in November 2025 (see Section 3.4). Source: BLS QCEW; Felten et al. (2021); Indeed Hiring Lab.

4 Empirical Framework

4.1 Core Continuous DiD Models

The primary empirical models are continuous Difference-in-Differences regressions. The core wage model is:

$$\log(\text{wage}_{it}) = \alpha_i + \delta_t + \beta (\text{Exposure}_i \times \text{AIshare}_t) + \varepsilon_{it} \quad (1)$$

where $\log(\text{wage}_{it})$ is the natural logarithm of average weekly wages in industry i in month t , α_i are industry fixed effects that absorb all time-invariant characteristics of each industry such as skill intensity, capital intensity, and historical wage levels, δ_t are monthly time fixed effects that absorb all macroeconomic shocks affecting all industries simultaneously such as business cycle fluctuations and aggregate price changes, and β is the parameter of interest capturing whether industries with higher pre-determined AI exposure experience differential wage changes as national AI-related hiring intensity rises. The interaction term $\text{Exposure}_i \times \text{AIshare}_t$ is the continuous treatment variable — it varies because the same national time-varying AI hiring shock is loaded differently by industries with different levels of structural AI exposure. A positive β would indicate that higher-exposure industries experienced faster wage growth as AI hiring rose nationally, consistent with a wage complementarity hypothesis. A null β — as found here — indicates no differential wage response across the exposure distribution during the sample period.

The employment model applies the same logic to log employment:

$$\log(\text{employment}_{it}) = \alpha_i + \delta_t + \beta (\text{Exposure}_i \times \text{AIshare}_t) + \varepsilon_{it} \quad (2)$$

where $\log(\text{employment}_{it})$ is the natural logarithm of quarterly average employment in industry i in quarter t . A positive aggregate β in the employment model would indicate that higher-exposure industries experienced faster employment growth as AI hiring rose — consistent with a complementarity hypothesis in which AI tools increase labour demand. A null aggregate β with a positive β within the low-exposure subgroup — as found here — is consistent with a reallocation hypothesis in which employment shifts toward industries at the lower end of the exposure distribution as AI-intensive sectors become more productive with fewer workers.

Both models are estimated on the balanced panel using the `feols()` function in R, which implements the within-group demeaning estimator. Standard errors are clustered at the industry level to account for serial correlation within industries over time. Robustness checks include two-way clustered standard errors at the industry and time level, and heterogeneity analysis by splitting industries into exposure terciles to examine whether the aggregate coefficient masks differential effects across the distribution.

4.2 Identification and Fixed Effects

Industry fixed effects absorb all time-invariant industry characteristics. Time fixed effects absorb all macroeconomic shocks common to all industries. The fixed effects are estimated using the within-group demeaning estimator in `feols()` in R. The identifying variation comes from a specific source: since AIshare_t varies only over time and Exposure_i varies only across industries, the interaction term $\text{Exposure}_i \times \text{AIshare}_t$ varies because the same national time shock is loaded differently by industries with different levels of AI exposure. This follows an exposure-based shift-share logic — industries with different pre-determined AI exposure are differentially loaded onto the same national AI-hiring shock — following the framework of Goldsmith-Pinkham et al. (2020). The coefficient β captures differential wage or employment changes conditional on all permanent industry differences and economy-wide time trends. Standard errors are clustered at the industry level.

4.3 Parallel Trends Testing

Parallel trends is assessed using two approaches. First, event study plots interacting industry exposure with time period indicators relative to ChatGPT launch are produced. The ChatGPT date is used as a reference not because it is the treatment in the main model but because it represents the most visible structural break in the AI-related hiring intensity series. A complementary binned event study groups Aishare variation into broad economic periods as a parallel trends diagnostic more directly connected to the continuous DiD. Second, formal pre-trend placebo tests assign a fake treatment date within the pre-period.

4.4 Methodological Notes

Construction and interpretation of the AIIE index

The AIIE index measures structural exposure to AI capabilities based on task content — not actual AI adoption. A high AIIE score means an industry employs workers whose tasks overlap with what current AI systems can do, not that the industry has invested more in AI. The coefficient β captures the differential response of structurally exposed industries to a national trend in AI-related hiring.

Dynamic effects and lagged wage responses

To examine delayed wage adjustment, the core wage model was re-estimated with lagged interaction terms at 3, 6, and 12 months. The results show no significant lagged response at any horizon — coefficients are -0.003 ($p = 0.283$) at three months, -0.002 ($p = 0.432$) at six months, and -0.001 ($p = 0.642$) at twelve months. Importantly the magnitude declines rather than rises with lag length, moving from -0.003 at three months to -0.001 at twelve months. If wages were simply adjusting with a delay we would expect the coefficient to grow larger at longer lags — the opposite pattern here is inconsistent with a simple delayed adjustment hypothesis and suggests that wage adjustment to AI-related hiring intensity operates over longer horizons than the current

six-year sample window captures.

Endogeneity of the AI posting share

Three features address potential endogeneity. First, the AI-related hiring intensity measure is a national aggregate identical across all industries. Second, the AIIE exposure index is pre-determined from occupational task characteristics. Third, fixed effects absorb permanent differences and aggregate trends. One residual concern is that high-exposure industries collectively may partly drive the national posting share — fully resolving this would require an external instrument, which is left for future research.

5 Results

5.1 Parallel Trends — Event Studies

Before presenting the main results, it is important to assess whether the parallel trends assumption holds. This step is not a formality: the entire empirical strategy rests on the claim that, absent the rise in national AI-related hiring intensity, high- and low-exposure industries would have followed similar wage and employment trajectories. If exposure groups were already diverging before AI-related hiring intensified, any divergence detected later could simply reflect the continuation of a pre-existing trend rather than a genuine response to AI. Figures 1 through 3 present event study plots for both outcomes. Pre-period coefficients — to the left of the dashed vertical line — should be small, close to zero, and show no systematic trend if parallel trends hold.




Figure 1: Event Study — AI Exposure Effect on Log Wage (Monthly). Each dot = coefficient on Exposure $\times$ Month relative to ChatGPT launch (reference = Month -1). Error bars = 95% CI. Sample: monthly wage panel 2019–2025, 145 industries. Pre-period coefficients are small and bounce around zero with no systematic trend, supporting parallel trends for wages.




Figure 2: Event Study — AI Exposure Effect on Log Wage (Binned Periods). Bins: $-3 = 2019$, $-2 = 2020$, $-1 = 2021\text{--}Q3\ 2022$, $1 = 2023$, $2 = 2024$, $3 = 2025$. Note: Q4 2022 is omitted as the reference period. Pre-period bins are close to zero with confidence intervals crossing zero, supporting parallel trends. This binned version directly tests parallel trends in the continuous DiD specification by grouping AIshare variation into broad economic periods.

Figure 1 shows the monthly wage event study. The pre-period coefficients spanning approximately 36 months before the ChatGPT launch are small with no systematic trend. This is important for the credibility of the wage results presented in Section 5.2: because there is no visible upward or downward drift in the pre-period, the null wage effect reported later cannot be attributed to exposure groups already converging or diverging for reasons unrelated to AI. Figure 2 provides the binned version — pre-period bins have coefficients close to zero with confidence intervals crossing zero. Because this binned specification groups time into the same broad periods used to construct the continuous AI-related hiring intensity measure, it speaks directly to whether the identifying variation exploited in the main wage model is contaminated by a pre-existing trend, rather than only to trends around the ChatGPT launch date specifically. The formal pre-trend placebo test yields -0.007 with $p = 0.059$ — not significant at five percent though marginal at ten percent. If this marginal result reflected a genuine pre-trend rather than sampling noise, it would imply that part of the null wage effect documented later could mask a small, pre-existing divergence between exposure groups; given that the coefficient is small in magnitude and only marginally significant,

however, it is treated here as a mild limitation rather than as evidence that invalidates the wage analysis, and is acknowledged as such throughout the remainder of the essay.




Figure 3: Event Study — AI Exposure Effect on Log Employment (Quarterly). Each dot = coefficient on Exposure $\times$ Quarter relative to ChatGPT launch (reference = Quarter -1). Error bars = 95% CI. Sample: quarterly employment panel 2019–2025, 125 industries. The pre-period shows some noise but no systematic directional trend. The formal pre-trend placebo test yields $\beta = -0.000$ ($p = 0.251$) — providing no evidence of a statistically significant differential pre-trend.

Figure 3 shows the quarterly employment event study. The pre-period shows some noise around quarters -10 to -12 but confidence intervals are wide and cross zero everywhere with no systematic trend. The formal pre-trend placebo test yields $\beta = -0.000$ with $p = 0.251$ — providing no evidence of a statistically significant differential pre-trend. This result is considerably cleaner than the wage placebo in Figure 2, and it is the reason that the employment findings in Section 5.3 — particularly the within-tercile employment gradient, the central novel result of this essay — are treated with somewhat greater confidence than the wage findings throughout the remainder of the discussion. In other words, the asymmetry in how well parallel trends holds across the two outcomes is itself informative: it suggests that whatever differential dynamics are detected later in employment are less likely to be artefacts of pre-existing divergence than any differential dynamics that might have been detected in wages.

5.2 Core Results: Wages

Table 3 presents the results of the core continuous DiD model for log weekly wages, directly addressing the first of this essay’s central questions: has rising national AI-related hiring intensity produced measurable wage divergence between industries with different structural exposure to AI? The main specification in column 1 yields a coefficient of -0.003 with a standard error of 0.002 and a p -value of 0.230 — not statistically significant at any conventional level. A one unit increase in the full-sample standardised AIIE index combined with a one percentage point rise in national AI-related hiring intensity is associated with a -0.3 percent change in wages — indistinguishable from zero.

Table 3: Core DiD — Effect of AI-Related Hiring Intensity on Log Weekly Wage

	(1) Main	(2) Two-way	(3) Low Exp	(4) Med Exp	(5) High Exp	(6) Pre-trend
Exposure $\times$ AI-Rel. Hiring	-0.003	-0.003	-0.009	-0.011	-0.001	
	(0.002)	(0.002)	(0.013)	(0.013)	(0.007)	
Exposure $\times$ Placebo Post						-0.007^*
						(0.004)
Observations	11,295	11,295	3,844	3,688	3,763	6,574
Adj. R^2	0.983	0.983	0.984	0.979	0.980	0.986
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Outcome = log average weekly wage. Sample: monthly 2019–2025, 145 industries, 11,295 observations.

Exposure = standardised AIIE (Felten et al., 2021); mean = 0, SD = 1. AI-related hiring intensity = national monthly US Indeed AI job posting share. Col 1: SEs clustered by industry. Col 2: two-way clustered. Cols 3–5: exposure tercile subsamples. Col 6: pre-trend placebo (placebo Post = 1 from July 2021). All models include industry and monthly time fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column 2 confirms robustness to two-way clustered standard errors, indicating that the null

result is not an artefact of underestimated standard errors arising from serial correlation in either the industry or the time dimension. Columns 3 through 5 report the tercile heterogeneity — all null across low (-0.009 , $p = 0.482$), medium (-0.011 , $p = 0.411$), and high exposure (-0.001 , $p = 0.936$). This uniformity across the exposure distribution is itself an informative result: it rules out the possibility that wage effects exist but are concentrated in one segment of the exposure distribution and are simply being masked in the aggregate specification. Column 6 reports the pre-trend placebo yielding -0.007 ($p = 0.059$) as discussed in Section 5.1. Taken as a whole, Table 3 indicates that wage adjustment to rising AI-related hiring intensity has not yet materialised anywhere along the exposure distribution over the 2019–2025 sample period. This is directly consistent with the cross-country evidence reviewed in Section 2.4 — Finnish, US, and Danish administrative data summarised by Kauhanen and Rouvinen (2025), Hartley et al. (2024), and Humlum and Vestergaard (2025) all report similarly null aggregate wage effects — and supports the view, developed further in the Discussion chapter, that wages are a slow-moving margin of adjustment relative to employment.

Figure 4 provides the visual complement showing average log weekly wages over time for high- and low-exposure groups. Both groups rise steadily in parallel with no visible change around the ChatGPT launch.




Figure 4: Average Log Weekly Wage by AI Exposure Group (2019–2025). Blue = above-median AIIE (high exposure); Orange = below-median AIIE (low exposure). Red shading = COVID-19 period. Dashed line = ChatGPT launch (November 2022). Both groups trend upward in parallel with no visible divergence after ChatGPT, consistent with the null wage result in Table 3.

The two lines in Figure 4 visually corroborate the regression result in Table 3 rather than merely repeating it in graphical form: even at the ChatGPT launch date — the most visible and widely recognised structural break in AI capabilities over the sample period — there is no discernible kink, level shift, or change in slope separating the high- and low-exposure groups. Given that this is precisely the date at which one might most plausibly expect a wage response to first appear if one existed, its absence here strengthens the conclusion that, at least within the time horizon studied, AI-related hiring intensity has not produced a detectable wage channel of adjustment for either highly exposed or lightly exposed industries.

5.3 Core Results: Employment

Table 4 presents the results of the continuous DiD model with log employment as the outcome, addressing the second central question of this essay: if wages have not adjusted, has employment adjusted instead, and if so, in which industries? The aggregate employment coefficient in column 1 is -0.000 with a p -value of 0.935 — clearly null. However the tercile heterogeneity analysis reveals the most significant and novel finding of this paper.

Table 4: Core DiD — Effect of AI-Related Hiring Intensity on Log Employment

	(1) Main	(2) Two-way	(3) Low Exp	(4) Med Exp	(5) High Exp	(6) Pre-trend
Exposure $\times$ AI-Rel. Hiring	−0.000	−0.000	0.038**	−0.013	−0.005	
	(0.004)	(0.005)	(0.015)	(0.025)	(0.012)	
Exposure $\times$ Placebo Post						−0.000
						(0.005)
Observations	3,125	3,125	1,050	1,050	1,025	1,750
Adj. R^2	0.995	0.995	0.997	0.986	0.997	0.994
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Outcome = log quarterly average employment. Sample: quarterly 2019–2025, 125 industries, 3,125

observations. Exposure = standardised AIIE (Felten et al., 2021); mean = 0, SD = 1. AI-related hiring intensity = quarterly average of monthly Indeed AI posting share. Col 1: SEs clustered by industry. Col 2: two-way clustered. Cols 3–5: exposure tercile subsamples. Col 6: pre-trend placebo (placebo Post = 1 during 2022). All models include industry and quarterly time fixed effects. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Column 3 shows that within the low-exposure tercile, a one unit increase in the full-sample standardised AIIE index is associated with a 3.8 percent differential employment gain for a one percentage point rise in national AI-related hiring intensity ($\beta = +0.038$, $p < 0.05$). This result survives two-way clustering and the pre-trend placebo test yields $\beta = -0.000$ ($p = 0.251$) — small and statistically insignificant. The medium-exposure coefficient is -0.013 ($p = 0.604$) and high-exposure is -0.005 ($p = 0.671$) — neither significant. The pattern across columns 3 through 5 is itself the central finding of this essay: rather than a uniform employment response, or an effect concentrated in the most AI-exposed industries as a simple displacement story might predict, the only statistically detectable employment response occurs within the *least* exposed segment of the distribution. Because the significant employment effect emerges from a heterogeneity specification rather than the primary aggregate model, it should be interpreted as suggestive rather than definitive

evidence of reallocation. Future work should test whether this pattern persists as additional post-2025 data become available.

This within-tercile exposure gradient is consistent with a labour reallocation hypothesis. As national AI-related hiring intensity rises, employment grows faster within the low-exposure group for industries that are relatively more AI-compatible even within that group — suggesting that even at the lower end of the AI exposure distribution, industries with slightly more AI-compatible tasks capture some reallocation benefits. Read against the displacement-reinstatement framework set out in Section 2.4, this result is consistent with a specific version of that mechanism: rather than reinstatement occurring within the most AI-exposed industries themselves, it appears to surface in industries positioned to absorb labour demand displaced from elsewhere in the economy, even when those industries are not themselves highly exposed to AI. This distinction matters for how the results should be read: the finding does not imply that AI is creating new AI-complementary jobs within high-exposure sectors, but rather that structurally low-exposure sectors with somewhat more AI-compatible task content appear to be where labour market adjustment is currently visible.




Figure 5: Average Log Employment by AI Exposure Group (2019–2025). Blue = above-median AIIE; Orange = below-median AIIE. Red shading = COVID-19 period. Dashed line = ChatGPT launch (Q4 2022). The 2019 pre-COVID baseline is visible — both groups follow similar pre-COVID trajectories before the pandemic disrupted low-exposure industries more severely.

Figure 5 shows that the two broad exposure groups — defined simply as above- versus

below-median AIIE — track each other closely before COVID and do not visibly diverge after the ChatGPT launch. This is an important complement to Table 4: it confirms that the significant employment result documented above is genuinely a *within*-tercile phenomenon, driven by variation in AIIE among low-exposure industries rather than by a broad divergence between the high- and low-exposure halves of the distribution as a whole. The COVID-19 disruption visible in the shaded region, and its apparently greater severity for the low-exposure group, foreshadows the descriptive COVID extension presented in Appendix A, but does not by itself drive the post-ChatGPT employment gradient documented in Table 4, since that gradient is estimated using the full continuous AIIE measure within the low-exposure subsample rather than the broad two-group split shown here.

To connect the regression results to concrete economic reality, Figure 6 presents raw employment changes at the individual industry level from Q1 2019 to Q3 2025. Three COVID-dominated leisure industries are excluded as their extreme post-pandemic rebounds are unrelated to AI.




Figure 6: Employment Change by Industry — Raw (2019 Q1 to 2025 Q3). Top 15 gaining and bottom 15 losing industries coloured by AI exposure tercile. Three COVID-dominated leisure industries excluded. Business Support Services (High exposure) is the single largest employment loser at approximately -35% . Low-exposure construction industries dominate the top gainers. Raw descriptive changes — do not control for industry or time fixed effects. See Table 4 for fixed-effects regression estimates.

Figure 6 shows that Business Support Services — a high-exposure industry performing administrative, scheduling, and document management tasks directly substitutable by current AI tools — is the single largest employment loser at approximately -35 percent. In contrast, three low-exposure construction industries appear among the top gainers. This contrast is precisely the one anticipated by the task-based discussion in the Introduction and the exposure-measurement literature reviewed in Section 2.3: Business Support Services sits at the cognitive, language-based,

information-processing end of the task spectrum that Felten et al. (2021) identify as most exposed to AI, while construction occupies the non-routine manual end of the spectrum that AI has so far shown little capacity to substitute for. Seeing this exact contrast emerge in the raw data — without yet controlling for any of the secular trends that could also explain it — provides an intuitive, descriptive anchor for the more rigorous regression evidence that follows. These raw changes reflect all factors over this period including COVID recovery and secular trends and should be interpreted descriptively. The decline in Business Support Services is consistent with an AI-substitution interpretation but cannot be attributed uniquely to AI because other industry-specific shocks — including post-COVID reorganisation of office work and outsourcing — may also have contributed.

Figure 7 provides a more rigorous version of the same analysis, controlling for industry and time fixed effects to isolate employment changes relative to common time trends. Positive values indicate employment grew faster than predicted by industry fixed effects and common quarterly time shocks after Q4 2022, while negative values indicate slower growth.




Figure 7: Employment Change After Controlling for Industry and Time Fixed Effects (Post vs Pre-ChatGPT). Employment residuals after partialling out industry fixed effects and common quarterly time shocks. Positive values = employment grew faster than predicted by fixed effects post-ChatGPT. Negative values = slower than predicted. Business Support Services (High exposure, -0.27) and Nondepository Credit Intermediation (High exposure, -0.17) are the two largest below-trend losers, consistent with AI substitutability of administrative and financial processing tasks. Above-trend gainers reflect ongoing post-COVID leisure recovery. Colours show AI exposure tercile based on AIIE index (Felten et al., 2021).

Figure 7 reveals that even after controlling for industry and time fixed effects, Business Support Services shows the largest post-ChatGPT decline at approximately -0.27 log points. Nondepository Credit Intermediation — another high-exposure financial services industry — is the second largest loser at -0.17 . The fact that the same industry, Business Support Services, remains the largest loser in both Figure 6 and Figure 7 is more informative than either figure alone:

Figure 6 could in principle be explained entirely by a secular decline that happened to coincide with the AI era, but Figure 7 removes exactly that possibility by partialling out each industry’s own average level and each quarter’s common shock. That the result survives this much stricter test makes the AI-substitution interpretation considerably more credible than a purely descriptive reading of the raw data would support on its own. The above-trend gainers include several leisure industries reflecting ongoing post-COVID recovery rather than an AI-related mechanism, which is an important caveat to bear in mind when interpreting the gainer side of this figure, even though the loser side is more directly informative for this essay’s purposes.

Figure 8 provides the most direct visual evidence for the main finding by splitting each exposure group into above- and below-median AIIE subgroups, directly visualising the within-tercile gradient identified in Table 4 column 3.




Figure 8: Employment Trends by Exposure Group and Within-Group AIIE Level (2019–2025). Five lines show mean log employment for subgroups defined by exposure tercile and within-tercile AIIE level. Red shading = COVID-19 period. Dashed line = ChatGPT launch (Q4 2022). Within the low-exposure group, industries with higher AIIE scores (dark orange) grow faster than those with lower AIIE scores (light orange), with the gap widening after 2021 and again after Q4 2022. This within-tercile divergence is consistent with the positive coefficient in Table 4 column 3 ($\beta = +0.038$, $p < 0.05$). Level differences across groups reflect permanent industry size differences absorbed by fixed effects in the regression.

Figure 8 shows that within the low-exposure group, the dark orange line consistently grows faster than the light orange line, with the gap widening after 2021 and again after the ChatGPT

launch. The timing of this widening is itself part of the evidence: a gap that was constant throughout the sample would be more consistent with a fixed structural difference between these two subgroups unrelated to AI, whereas a gap that widens specifically as national AI-related hiring intensity rises, and again after the most visible AI-specific shock in the sample period, is more consistent with the gradient being driven by AI-related hiring intensity itself rather than by some other time-invariant industry characteristic. This figure is, in effect, the within-tercile coefficient from Table 4 made visible: the two orange lines are the same low-exposure industries used to estimate column 3, simply split at their own median AIIE score rather than pooled together.

Figure 9 focuses on the low-exposure tercile only, confirming the positive within-group slope, while Figure 10 shows the three-panel scatter across all terciles.




Figure 9: Employment Change vs AI Exposure — Low Exposure Industries Only (2019–2025). 42 industries in the bottom exposure tercile. Red line = OLS fit with 95% CI. The positive slope within the low-exposure group is consistent with the within-tercile gradient in Table 4 column 3 ($\beta = +0.038$, $p < 0.05$).




Figure 10: Employment Change vs AI Exposure by Tercile (2019–2025). Each panel = one exposure tercile. Each dot = one 4-digit NAICS industry. Red line = OLS fit with 95% CI. Positive slope in Low Exposure panel is consistent with the within-tercile gradient. Flat slope in Medium and negative slope in High are consistent with the null and negative (though insignificant) coefficients in Table 4 columns 4 and 5.

Figures 9 and 10 represent the same underlying employment-change variable used to construct Table 4, viewed at the level of individual industries rather than as a single pooled coefficient. Their value lies less in introducing new information than in demonstrating that the regression result is not an artefact of a particular functional form or estimation choice: the positive association in the low-exposure tercile is visible directly in the scatter of individual industries, not only in a fixed-effects regression coefficient. The progression from a positive slope in the low-exposure panel, to an essentially flat slope in the medium-exposure panel, to a negative (though statistically insignificant) slope in the high-exposure panel in Figure 10 traces out the full heterogeneity pattern reported numerically in Table 4 columns 3 through 5, and is consistent with the asymmetric reallocation mechanism discussed above: the gradient is positive only at the bottom of the exposure distribution, flattens through the middle, and turns mildly negative — without reaching significance — at the top, where the competing displacement and productivity effects discussed in Section 2.2 are most likely to be operating simultaneously.

5.4 Supplementary: ChatGPT Binary DiD

Table 5 presents the supplementary ChatGPT binary DiD results, which serve as a robustness check on the continuous DiD results above using an entirely different definition of treatment: rather than relying on the continuously varying national AI posting share, this specification treats the ChatGPT launch as a single discrete event. If the within-tercile employment gradient documented in Table 4 reflects a genuine AI-related phenomenon rather than an artefact of the particular continuous treatment measure used, it should also appear, and ideally strengthen, when generative AI's most visible and widely recognised structural break is used as the treatment definition instead. These results should be interpreted with caution given the short post-treatment period of approximately two and a half years. Although ChatGPT launched in November 2022, Q4 2022 is excluded from the balanced employment panel due to insufficient industry coverage in the QCEW files for that quarter. The first observed post-launch quarter in the panel is therefore Q1 2023, and the Post-ChatGPT indicator equals one from Q1 2023 onwards. Given the short post-period these results should be treated as suggestive rather than causal, consistent with the caution advised by Angrist and Pischke (2009) regarding inference in short panels.

Table 5: Supplementary — ChatGPT Binary DiD — Wages and Employment

	(1) Wage All	(2) Wage Low	(3) Wage High	(4) Empl All	(5) Empl Low	(6) Empl High
Exposure $\times$ Post-ChatGPT	−0.005 (0.005)	−0.031 (0.022)	0.021 (0.020)	−0.002 (0.008)	0.058** (0.028)	−0.034 (0.025)
Observations	3,125	1,050	1,025	3,125	1,050	1,025
Adj. R^2	0.983	0.985	0.980	0.995	0.997	0.997
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: SUPPLEMENTARY ANALYSIS — interpret with caution. Post-ChatGPT = 1 from Q1 2023 (first observed post-launch quarter in the balanced employment panel; Q4 2022 excluded due to insufficient industry coverage in QCEW files). Cols 1–3: log wage outcome. Cols 4–6: log employment outcome. Low/High = bottom and top exposure tercile subsamples. Short post-period (2023–2025) limits causal interpretation. Industry and quarterly time fixed effects.

SEs clustered by industry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

All wage coefficients are null, reproducing the pattern from Table 3 under a different treatment definition and reinforcing the conclusion that wages have not yet responded to AI-related hiring intensity by any reasonable measure of the AI shock considered in this essay. For employment, column 5 shows that within the low-exposure tercile, a one unit increase in the full-sample standardised AIIE index is associated with a 5.8 percent differential employment gain after ChatGPT ($\beta = +0.058$, $p < 0.05$). High-exposure industries show a negative but insignificant coefficient ($\beta = -0.034$, $p = 0.181$), directionally consistent with the reallocation hypothesis. The same within-tercile gradient appears in two different specifications — continuous DiD with AIshare and binary DiD with ChatGPT Post — which provides additional consistency across different identification approaches: it is considerably less likely that two different treatment definitions, constructed from different aspects of the same underlying phenomenon, would coincidentally produce the same significant subgroup pattern by chance. The increase in magnitude from $\beta = +0.038$ in the continuous specification to $\beta = +0.058$ here is suggestive of an intensification of the reallocation dynamic specifically in the post-ChatGPT period, consistent with the idea that

the generative AI era represented a step change relative to the more gradual pre-ChatGPT rise in AI-related hiring intensity. This interpretation should, however, be offered cautiously: with only the short post-ChatGPT period, the larger coefficient could equally reflect greater estimation noise in a shorter, more recent sample rather than a genuine acceleration, and this essay does not attempt to adjudicate between these two explanations definitively.

6 Discussion

Together, the results are consistent with a coherent picture of how national AI-related hiring intensity is reshaping US industry labour markets. The evidence consistently points to employment rather than wages as the primary channel through which AI exposure manifests in the data, and the patterns are concentrated within the low-exposure group rather than in the high-exposure industries most directly threatened by AI technology.

The null wage result is consistent with the view that wage adjustment to technology shocks is slow and sticky. The finding aligns with Acemoglu and Restrepo’s (2018) theoretical prediction that automation’s effects appear in employment quantities before translating into wage changes. It is also consistent with independent evidence from Finland (Kauhanen and Rouvinen, 2025), the United States (Hartley et al., 2024), Denmark (Humlum and Vestergaard, 2025), and Anthropic’s own usage data (Massenkoff and McCrory, 2026) — four entirely different data sources converging on the same null wage result.

The significant positive within-tercile employment gradient in low-exposure industries is the most novel finding. Figures 6, 7, and 8 connect the abstract regression results to economic reality. Figure 6 shows in raw terms that Business Support Services lost approximately 35 percent of employment while low-exposure construction industries were among the top gainers. Figure 7 confirms that Business Support Services and Nondepository Credit Intermediation show the largest post-ChatGPT declines after controlling for industry and time fixed effects. Figure 8 directly visualises the within-tercile gradient — the gap between dark and light orange lines widens as AI hiring accelerates. The decline in Business Support Services is consistent with an AI-substitution interpretation, though other industry-specific factors may also have contributed and this should not be attributed uniquely to AI.

A natural question is why the within-tercile gradient appears specifically in the low-exposure group and not in the high-exposure group. One explanation is that the reallocation mechanism operates asymmetrically. High-exposure industries face two competing forces simultaneously — AI substitution reduces demand for labour on the one hand, while AI-driven productivity increases may raise output and therefore labour demand on the other. These forces partially offset each other, producing a null net employment effect in the high-exposure group. In contrast, low-exposure industries face primarily the reallocation benefit without the direct substitution effect. Within the low-exposure group, industries with slightly higher AIIE scores are better positioned to absorb reallocated labour because their workers have somewhat more AI-compatible skills. This interpretation is consistent with Johnston and Makridis (2025), who find positive employment effects specifically where AI requires human collaboration.

The supplementary ChatGPT analysis provides suggestive evidence that these patterns intensified after November 2022 — the within-tercile gradient becomes larger ($\beta = +0.058$ vs $+0.038$) and remains significant. The short post-period means this should be treated as a hypothesis for future investigation. A COVID extension presented in Appendix A finds descriptive evidence of differential employment patterns during the pandemic but the result is not robust to controlling for remote work capacity (Dingel and Neiman, 2020).

7 Conclusion

This paper estimates the association between national AI-related hiring intensity and wages and employment across US industries using a continuous Difference-in-Differences framework. Three main findings emerge. First, rising AI-related hiring intensity is not associated with statistically significant differential wage changes across industries with different levels of AI exposure — consistent with the view that wage adjustment operates over longer horizons and with independent evidence from four different datasets using completely different methods. Second, within the low-exposure industry group, a one unit increase in the full-sample standardised AIIE index is associated with a 3.8 percent differential employment gain for a one percentage point increase in national AI-related hiring intensity ($\beta = +0.038$, $p < 0.05$), consistent with a labour reallocation hypothesis. Third, a supplementary ChatGPT binary DiD confirms this within-tercile employment pattern ($\beta = +0.058$, $p < 0.05$), with the pattern visible both in regression results and in industry-level descriptive figures showing Business Support Services as the largest employment loser and low-exposure construction industries as top gainers. Because these results emerge from a heterogeneity specification and a short post-treatment window, they should be interpreted as suggestive evidence of reallocation patterns rather than definitive causal estimates.

These findings have several policy implications. First, policies focused solely on retraining workers for high AI-exposure roles may miss an important part of the adjustment process — the evidence suggests that employment is also growing within the low-exposure group among industries with slightly more AI-compatible tasks. Supporting workers in these industries through skills development and occupational mobility programmes may be as important as preparing workers for high-AI roles. Second, the absence of wage effects through 2025 suggests that the distributional consequences of AI-related hiring intensity have not yet materialised in pay — but this does not mean they will not. Monitoring wage trends in both high and low exposure industries as the

post-ChatGPT window extends will be essential for identifying early signs of divergence before they become entrenched.

Future research should extend this analysis in several directions. As longer post-ChatGPT data become available the supplementary analysis will become substantially more informative — the current three-year post-period is too short to draw strong causal conclusions about the generative AI era. Industry-specific AI adoption measures — such as those from Lightcast or the Stanford AI Index — would provide richer treatment variation and allow separation of AI structural exposure from actual adoption decisions. Occupation-level analysis using matched employer-employee data would allow more precise estimation of distributional effects across workers with different skill profiles within the same industry. Finally extending the geographic dimension to exploit variation in AI research concentration across US commuting zones — following the approach of Acemoglu and Restrepo (2020) for robots — would provide a stronger identification strategy than the industry-level shift-share design used here.

References

- Acemoglu, D., & Autor, D. H. (2011). Skills, tasks and technologies: Implications for employment and earnings. In O. Ashenfelter & D. Card (Eds.), *Handbook of Labor Economics* (Vol. 4B, pp. 1043–1171). Elsevier.
- Acemoglu, D., & Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6), 1488–1542.
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from US labor markets. *Journal of Political Economy*, 128(6), 2188–2244.
- Angrist, J. D., & Pischke, J. S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press.
- Autor, D. H. (2019). Work of the past, work of the future. *AEA Papers and Proceedings*, 109, 1–32.
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333.
- Brynjolfsson, E., Chandar, B., & Chen, R. (2025). Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence. *Stanford Digital Economy Lab Working Paper*. <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>
- Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–941.

- Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108, 43–47.
- Bureau of Labor Statistics. (2025). *Quarterly Census of Employment and Wages (QCEW)*. U.S. Department of Labor. <https://www.bls.gov/cew/>
- Deming, D. J. (2017). The growing importance of social skills in the labor market. *Quarterly Journal of Economics*, 132(4), 1593–1640.
- Dingel, J. I., & Neiman, B. (2020). How many jobs can be done at home? *Journal of Public Economics*, 189, 104235.
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). GPTs are GPTs: An early look at the labor market impact potential of large language models. *Science*, 381(6654), 155–158.
- Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217.
- Frey, C. B., & Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114, 254–280.
- Goldsmith-Pinkham, P., Sorkin, I., & Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*, 110(8), 2586–2624.
- Goos, M., & Manning, A. (2007). Lousy and lovely jobs: The rising polarization of work in Britain. *Review of Economics and Statistics*, 89(1), 118–133.
- Graetz, G., & Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5), 753–768.
- Hartley, J., Jolevski, F., Melo, V., & Moore, B. (2024). The labor market effects of generative artificial intelligence. *SSRN Working Paper 5136877*.

- Humlum, A., & Vestergaard, E. (2025). Large language models, small labor market effects. *NBER Working Paper No. 33777*.
- Hui, X., Reshef, O., & Zhou, L. (2024). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science*, 35(6), 1977–1989.
- Indeed Hiring Lab. (2024). *AI job postings tracker — Indeed*. <https://github.com/hiring-lab/indeed-ai-tracker>. Creative Commons Attribution 4.0 International License.
- Johnston, A., & Makridis, C. (2025). AI, output, and employment. *SSRN Working Paper 5375017*.
- Kauhanen, A., & Rouvinen, P. (2025). Assessing early labour market effects of generative AI: Evidence from population data. *Applied Economics Letters*.
- Massenkoff, M., & McCrory, P. (2026). Labor market impacts of AI: A new measure and early evidence. *Anthropic*.
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192.
- PwC. (2025). *Global AI jobs barometer 2025*. PricewaterhouseCoopers. <https://www.pwc.com/gx/en/services/ai/ai-jobs-barometer.html>
- Webb, M. (2020). The impact of artificial intelligence on the labor market. *Stanford University Working Paper*.

A COVID Extension

This appendix presents the COVID extension analysis as a descriptive exercise. The model is:

$$\log(\text{employment}_{it}) = \alpha_i + \delta_t + \beta (\text{Exposure}_i \times \text{COVID}_t) + \varepsilon_{it} \quad (\text{A1})$$

where COVID_t equals one from Q1 2020 to Q2 2021. The employment panel was extended to include 2019 as a pre-COVID baseline. Table A1 presents the baseline results without a remote work control. Following the concern raised by the supervisor and consistent with the literature (Dingel and Neiman, 2020), Table A2 adds an interaction between each industry's teleworkable employment share and the COVID indicator to separate AI exposure effects from remote work capacity effects.

Table A1: COVID Extension — Effect of AI Exposure on Wages and Employment

	(1) Wage	(2) Empl	(3) Two-way	(4) Low Exp	(5) Med Exp	(6) High Exp	(7) Pre-trend
Exposure $\times$ COVID	0.004 (0.004)	0.015* (0.008)	0.015 (0.009)	−0.073** (0.030)	0.078 (0.066)	0.027 (0.022)	
Exposure $\times$ Placebo							−0.006** (0.002)
Observations	3,125	3,125	3,125	1,050	1,050	1,025	500
Adj. R^2	0.983	0.995	0.995	0.997	0.986	0.997	0.999
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: COVID = 1 from Q1 2020 to Q2 2021. Pre-COVID baseline = Q1 2019 to Q4 2019 (4 quarters). Col 1: log wage outcome. Cols 2–7: log employment outcome. Cols 4–6: exposure tercile subsamples. Col 3: two-way clustered SE. Col 7: pre-trend placebo (fake COVID = Q3–Q4 2019). All models include industry and quarterly time fixed effects. NOTE: The pre-trend placebo is significant (col 7: -0.006 , $p < 0.05$), indicating pre-existing employment divergence before COVID. Results should be interpreted descriptively rather than causally. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A1 shows that without controls, higher AI exposure is associated with stronger employment outcomes during COVID ($\beta = +0.015$, $p < 0.10$) and low-exposure industries suffered significant losses ($\beta = -0.073$, $p < 0.05$). However the pre-trend placebo in column 7 yields a significant coefficient (-0.006 , $p < 0.05$), indicating pre-existing employment divergence. This undermines the causal DiD interpretation — these results should be read as descriptive associations rather than causal estimates.

Table A2: COVID Extension with Remote Work Control (Dingel and Neiman, 2020)

	(1) Wage	(2) No ctrl	(3) Remote	(4) Two-way	(5) Low Exp	(6) Med Exp	(7) High Exp
Exposure $\times$ COVID	0.004 (0.004)	0.015* (0.008)	0.020 (0.012)	0.020 (0.012)	-0.078** (0.036)	0.122** (0.059)	0.045 (0.040)
Remote Work $\times$ COVID			-0.007 (0.013)	-0.007 (0.012)	0.074* (0.040)	-0.083** (0.037)	-0.011 (0.016)
Observations	3,125	3,125	3,125	3,125	1,050	1,050	1,025
Adj. R^2	0.983	0.995	0.995	0.995	0.997	0.986	0.997
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Model A (col 2): no remote work control. Model B (cols 3–7): includes Remote Work $\times$ COVID interaction.

Remote work index = share of teleworkable employment from Dingel and Neiman (2020), matched to 4-digit NAICS using 3-digit NAICS concordance. Both indices standardised to mean = 0, SD = 1. All models include industry and quarterly time fixed effects. SEs clustered by industry except col 4 (two-way clustered). * $p < 0.10$, ** $p < 0.05$,

*** $p < 0.01$.

Table A2 shows that after controlling for remote work capacity, the aggregate Exposure $\times$ COVID coefficient loses statistical significance ($\beta = +0.020$, $p > 0.10$). The Remote Work $\times$ COVID coefficient is also insignificant in the aggregate model. The heterogeneity analysis reveals nuanced patterns — low-exposure industries show significant employment losses after the remote work control (-0.078 , $p < 0.05$) while medium-exposure industries show significant gains ($+0.122$, $p < 0.05$). Overall the COVID extension is presented as a descriptive exercise — the parallel trends assumption is not fully supported and the aggregate result is not robust to controlling for remote work capacity.




Figure A1: COVID Heterogeneity — Employment Effect by AI Exposure Tercile. Bars = coefficient on Exposure $\times$ COVID from Table A1. Error bars = 95% CI. Pre-COVID baseline = 2019 (4 quarters). Low-exposure industries show a significant negative employment effect (-0.073 , $p < 0.05$). Medium and high exposure show positive but insignificant effects. Results should be treated as descriptive given the failed parallel trends test.