Generative AI and the labour market: a task-based empirical approach

Matteo Caruso[a]

[a]*Department of Economics, Social Studies, Applied Mathematics and Statistics (ESOMAS), University of Turin, Italy*

Abstract

This paper evaluates the early labour-market impact of Generative Artificial Intelligence (AI) on the United States using the public release of ChatGPT in November 2022 as a natural experiment. I combine IPUMS ACS micro-data (2021–2024) with a novel, revealed-preference measure of occupational AI exposure built from the Anthropic Economic Index (AEI, Handa et al. 2025), which classifies millions of real Claude conversations into O*NET task categories and distinguishes between *automation* (full delegation) and *augmentation* (iterative use) interactions. Starting from this task-level evidence, I aggregate exposure at the occupation level and group the 974 Job-Explorer occupations into three mutually exclusive categories — Automation-dominant, Augmentation-dominant, and Irrelevant — used as the treatment structure in a difference-in-differences design. I find no significant effect of ChatGPT exposure on aggregate employment, a positive wage effect of around three log-points for Automation-dominant occupations, and an increase in usual weekly hours of roughly 0.5–0.9 hours in both exposed groups. Women gain disproportionately from the augmentation channel across wages, employment, and hours. The results are consistent with an interpretation in which workers adopted ChatGPT privately to accelerate their most automatable tasks and translated the time saved into higher productivity rather than displacement.

Keywords: Artificial Intelligence, labour market, skill-biased, difference-in-difference, employment

---

*Please address correspondence to Matteo Caruso

Email address: `m.caruso@unito.it` (Matteo Caruso)

Introduction

The public release of ChatGPT on 30 November 2022 is arguably the fastest diffusion of a consumer technology on record. By mid-2025, the tool had reached over 700 million weekly active users and was receiving on the order of 18 billion messages per week, with more than two thirds of these interactions involving work- or productivity-related tasks (Chatterji et al., 2025). Bick et al. (2025) show that, less than two years after launch, around 40% of U.S. working-age adults had already used Generative AI, a diffusion rate much faster than that of personal computers or the internet. The speed and scale of this diffusion raise a natural question: what has this early wave of Generative AI meant for the U.S. labour market?

A well-established literature, building on the task-based framework of Autor et al. (2003) and Acemoglu and Autor (2011), argues that the labour-market footprint of any new technology depends on *which tasks* within each occupation it can automate, augment, or leave untouched. For Generative AI, a rich set of "exposure indices" has been produced from expert ratings or LLM self-assessments of O*NET tasks (Eloundou et al., 2024; Felten et al., 2023; Gmyrek et al., 2025b). These measures, however, remain *potential*: they tell us what LLMs *could* do, not what users actually *ask them to do*. The Anthropic Economic Index (Handa et al., 2025) fills this gap by clustering millions of real Claude conversations and mapping them onto O*NET tasks, distinguishing between automation (full delegation) and augmentation (iterative, feedback-loop) interactions.

This paper puts the AEI to work in a causal design. I aggregate the task-level Claude data at the occupation level and classify the 974 Job-Explorer occupations into three mutually exclusive groups: *Automation-dominant* (above-median exposure with automation shares exceeding augmentation shares), *Augmentation-dominant* (above-median exposure with augmentation prevailing), and *Irrelevant* (below-median exposure). I merge these groups into the IPUMS ACS 2021–2024 and estimate a difference-in-differences specification that interacts each group with a post-ChatGPT indicator, controlling for year, state, and occupation fixed effects, standard demographic controls, and person weights.

Three main findings emerge. First, there is no detectable effect of ChatGPT exposure on the probability of employment, in line with the null diffusion effects documented by Humlum and Vestergaard (2024), Cazzaniga et al. (2024), and the papers surveyed in Guarascio and Reljic (2025). Second, wages in Automation-dominant occupations rise by about three log-points after 2022 relative to the Irrelevant group, while no comparable effect appears for Augmentation-dominant occupations — an asymmetry that is consistent with early ChatGPT being better suited to one-shot delegation than to iterative use, and with the productivity gains documented by Brynjolfsson et al. (2023) and Czarnitzki et al. (2023). Third, usual weekly hours rise in both exposed groups (by 0.5–0.9 hours per week), echoing survey-based evidence on AI-induced work intensification (Microsoft and LinkedIn, 2024). The heterogeneity analysis further shows that women are the main beneficiaries of the augmentation channel, with positive and significant triple interactions on wages and employment.

These results matter because they bear directly on the policy narrative surrounding Generative AI. Contrary to the displacement-centred view that dominated the early public debate, the first two years of ChatGPT appear to have produced a mostly *positive* shock for the workers most directly exposed, at least in the short run. The paper makes three contributions to the literature. It is, to my knowledge, the first to use the AEI as a *pre-treatment* measure of task-level AI exposure inside a natural-experiment design; it separates the automation and augmentation channels of Acemoglu and Autor (2011) into distinct, mutually exclusive treatment groups; and it explicitly confronts the identification challenge raised by the fact that the AEI is built on Claude rather than ChatGPT, via the assumption that users do not meaningfully distinguish between LLMs at the task level. Direct ChatGPT-usage data of the kind assembled by Chatterji et al. (2025) would further strengthen this mapping and is a natural extension of the present work.

The remainder of the paper is organised as follows. Section 1 reviews the relevant literature on task-based technological change, Generative-AI productivity, and AI-exposure measurement. Section 2 introduces the data and reports descriptive statistics, including a monthly IPUMS-CPS plot of employment by exposure group. Section 3 presents the DID and heterogeneity estimates. Section 4 concludes and discusses the main caveats.

1. Literature Review

1.1. Technology, tasks, and the labour market

A long tradition in labour economics studies how technological change reshapes the demand for skills and the allocation of workers across tasks. Autor et al. (1998) document that the diffusion of personal computers raised the relative demand for college-educated workers and contributed to the rise in the college wage premium observed since the 1980s. Autor et al. (2003) formalise this intuition in the canonical *task-based* framework: jobs are bundles of routine and non-routine tasks, and computer capital is a near-perfect substitute for codifiable, routine cognitive tasks and a complement to non-routine analytical and interactive ones. The framework rationalises the polarisation of employment (Autor and Dorn, 2013; Autor et al., 2006) and, more generally, the view that the labour-market impact of a new technology depends less on *which occupations* exist and more on *which tasks* within those occupations can be automated or augmented.

Acemoglu and Autor (2011) extend this reasoning into a unified supply-and-demand model in which technological change biased toward certain tasks has distributional effects mediated by the elasticity of substitution between tasks and factors. Acemoglu et al. (2022) apply the same logic to early Artificial Intelligence (AI), showing that AI-exposed establishments posted fewer non-AI vacancies but did not display measurable aggregate employment or wage effects, consistent with the idea that the labour-market footprint of any individual technology is an empirical question that must be answered on the *task* margin. This paper adopts the same task-level lens, but focuses on a distinct technological wave: the consumer-facing diffusion of Large Language Models (LLMs) after November 2022.

1.2. Generative AI, productivity, and employment

A growing empirical literature documents that Generative AI affects output and employment at the firm and worker level. Czarnitzki et al. (2023) exploit German firm-level innovation data and find that AI adoption raises labour productivity, especially for firms with complementary digital capabilities. Brynjolfsson et al. (2023) study the rollout of a Generative-AI customer-support assistant and show large gains in agent productivity (around 14%), concentrated among the least experienced workers, suggesting a compression of the within-firm productivity distribution. At the industry level, Guarascio and Reljic (2025) use European data to document heterogeneous employment and wage effects of AI adoption across sectors, with generally positive employment responses in AI-intensive activities.

In parallel, a set of recent papers use survey data to measure adoption at the individual level. Bick et al. (2025) show that, by mid-2024, roughly 40% of U.S. working-age adults had used Generative AI, making its diffusion faster than that of personal computers or the internet. Humlum and Vestergaard (2024) use Danish register data linked to a survey on ChatGPT and find widespread use but, so far, small effects on earnings and working hours — a finding echoed in the macro-micro projections of Cazzaniga et al. (2024) and in the early impact estimates of Crane and Soto (2026). Together, these studies point to a technology whose *adoption* is unusually rapid but whose *labour-market consequences* remain uncertain and plausibly mediated by the task composition of each occupation.

1.3. Measuring AI exposure: from theoretical scores to revealed use

Because Generative AI is difficult to observe directly in standard labour statistics, a large methodological literature builds indices of *AI exposure* at the occupation or task level. Felten et al. (2023) map LLM capabilities to occupational abilities; Eloundou et al. (2024) assess the share of tasks in each O*NET occupation that could be at least partially performed by GPT-4 and conclude that around 80% of U.S. workers are exposed in some of their tasks. Zarifhonarvar (2024) produces an analogous classification with a focus on white-collar jobs, and Aldasoro et al. (2024) specialise the exercise to finance. At the global level, Gmyrek et al. (2025b) and Gmyrek et al. (2025a) refine the ILO index of occupational exposure and translate it into projected employment effects, while Haque and Li (2025) documents heterogeneity in exposure across demographic groups.

A common limitation of these indices is that they are *potential* measures: they rely on expert ratings or LLM self-assessments of what could be automated, not on what is actually done with AI. Handa et al. (2025) address this gap by building the Anthropic Economic Index (AEI) from millions of real Claude conversations clustered with the Clio semantic-similarity pipeline and mapped onto O*NET tasks; for each task they estimate the share of interactions in which the human uses AI in an *automation* (directive) mode versus an *augmentation* (iterative, feedback-loop) mode. Handa and co-authors thus move the literature from "what jobs *could* be affected" to "what tasks are *actually* performed with AI", providing a revealed-preference measure that is, to my knowledge, the most granular available empirical benchmark of Generative-AI use.

1.4. *Contribution and identification strategy*

This paper fills three related gaps. First, most of the exposure literature stops at cross-sectional characterisations or static projections; I use the AEI as a *pre-treatment* measure of task-level AI relevance and embed it into a difference-in-differences design around the release of ChatGPT (November 2022), following the intuition in the productivity literature (Czarnitzki et al., 2023; Brynjolfsson et al., 2023; Guarascio and Reljic, 2025) but with representative household data. Second, by exploiting the mutually exclusive classification of occupations into *automation-dominant*, *augmentation-dominant*, and *irrelevant* groups, I separate the two channels that Acemoglu and Autor (2011) and the task-based tradition identify as conceptually distinct but empirically entangled. Third, I address the identification concern that the AEI is Claude-specific by assuming that end-users do not meaningfully distinguish between LLMs at the task level: the same tasks that Claude users offload to AI are, plausibly, those that ChatGPT users also offload, so that AEI-based exposure is informative about the treatment intensity induced by ChatGPT's release. Together, these choices deliver a revealed-preference, event-based evaluation of Generative AI in the U.S. labour market.

2. Data and descriptive statistics

The analysis merges IPUMS ACS micro-data (2021–2024, ages 16–64) with occupation-level exposure measures built from Anthropic's Job Explorer data (Handa et al. 2025). For each of the 974 O*NET occupations in the Job Explorer, I compute: (i) a *total exposure* intensity $\ln(1 + \text{conversations})$, (ii) an *automation share* and an *augmentation share* obtained by weighting each task's automation/augmentation percentage by its conversation count. From these I construct three mutually exclusive occupation dummies: *Irrelevant* (below-median total exposure, baseline control); *Automation-dominant* (at/above-median exposure with automation $>$ augmentation); and *Augmentation-dominant* (at/above-median exposure with augmentation $\geq$ automation).

Table 1 reports sample means. Table 3 compares the three outcomes (employment, log wage, usual weekly hours) across the three exposure groups. Tables 4–5 list the most represented occupations and NAICS-2 industries.

Table 1: Summary statistics, IPUMS ACS 2021–2024

	Mean	SD
Employed (=1)	0.900	0.300
Log wage income	10.850	0.779
Usual hours worked per week	42.765	9.763
Age	44.109	11.925
Female (=1)	0.367	0.482
College degree or above (=1)	0.226	0.418
Married (=1)	0.610	0.488
White (=1)	0.654	0.476
Observations	123913	

Table 2: IPUMS ACS sample split by AI-exposure group, 2021–2024

Group	Observations	Weighted persons	Share (%)
Augmentation	4,869,310	4,869,310	38.03
Automation	4,348,815	4,348,815	33.97
Irrelevant	3,585,330	3,585,330	28.00

Table 4: Top 15 occupations by IPUMS person-weight share

Occupation	Group	Share (%)
First-Line Supervisors of Office and Administrative Support Workers	Augmentation	38.03
First-Line Supervisors of Production and Operating Workers	Irrelevant	25.80
First-Line Supervisors of Construction Trades and Extraction Workers	Automation	21.72
First-Line Supervisors of Mechanics, Installers, and Repairers	Automation	8.19
Weighers, Measurers, Checkers, and Samplers, Recordkeeping	Automation	2.75
First-Line Supervisors of Agricultural Crop and Horticultural Workers	Irrelevant	1.49
Bioinformatics Technicians	Automation	1.30
Forest and Conservation Workers	Irrelevant	0.71

Table 3: Outcomes by exposure group

	Mean	SD
Augmentation		
Employed (=1)	0.900	0.301
Log wage income	10.776	0.757
Usual hours worked per week	40.559	8.806
Automation		
Employed (=1)	0.893	0.310
Log wage income	10.912	0.829
Usual hours worked per week	44.105	10.632
Irrelevant		
Employed (=1)	0.910	0.286
Log wage income	10.878	0.739
Usual hours worked per week	44.111	9.341
Total		
Employed (=1)	0.900	0.300
Log wage income	10.850	0.779
Usual hours worked per week	42.765	9.763
Observations	123913	

Table 5: Top 15 industries by IPUMS person-weight share

Industry	Share (%)
Manufacturing	25.31
Construction	21.58
Health Care and Social Assistance	7.07
Public Administration	6.96
Professional, Scientific, and Technical Services	5.32
Retail Trade	4.85
Finance and Insurance	4.17
Administrative and Support and Waste Management	3.97
Transportation and Warehousing	3.08
Utilities	3.03
Other Services (except Public Administration)	2.44
Agriculture, Forestry, Fishing and Hunting	1.93
Educational Services	1.83
Mining, Quarrying, and Oil and Gas Extraction	1.74
Real Estate and Rental and Leasing	1.73

Figure 1 plots monthly employment rates from the IPUMS-CPS Basic Monthly samples

(2022–2023) for each of the three exposure groups, indexed to January 2022. The vertical dotted line marks November 2022.

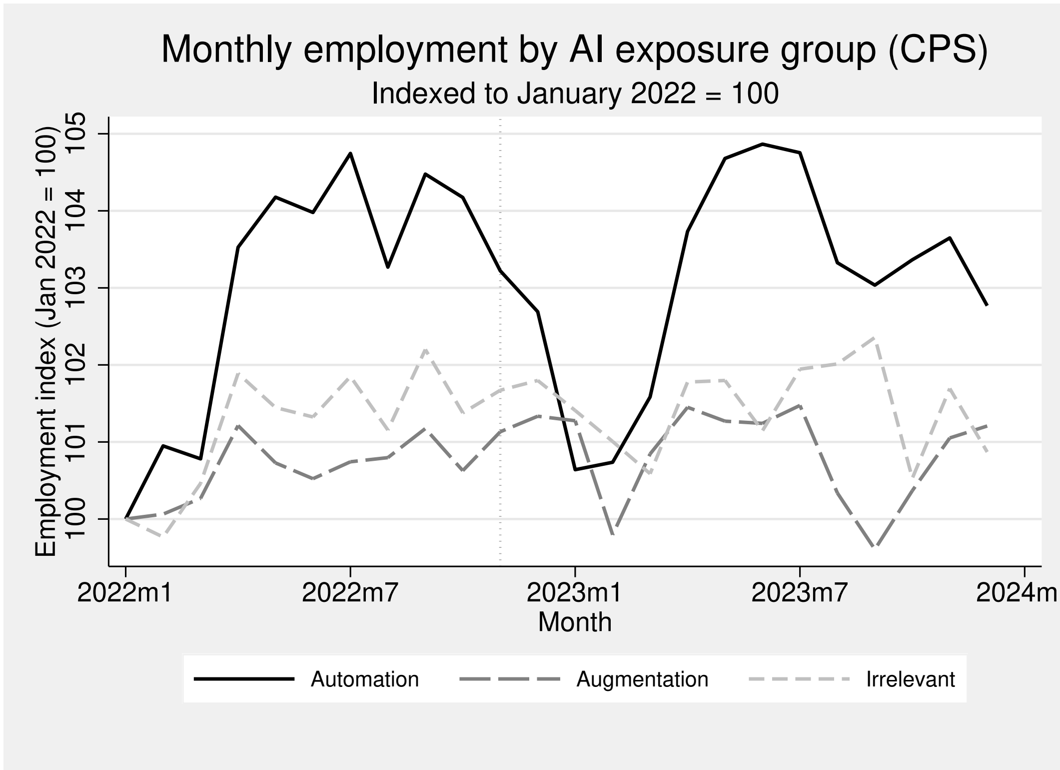

Figure 1: Monthly employment rate by AI exposure group, IPUMS-CPS 2021–2024 (Jan 2022 = 100).

A notable feature of Figure 1 is that the Automation-dominant series exhibits the largest month-to-month variability, markedly larger than the Augmentation and Irrelevant series. This higher baseline volatility is consistent with the Automation-dominant group being the most cyclically sensitive, and also suggests that, on prior grounds, I should expect the strongest post-ChatGPT reaction to come from this group — as indeed turns out to be the case in the wage and hours regressions below.

3. Results

Treatment is $\text{Post}_t = \mathbf{1}\{\text{year} \geq 2023\}$, interacted with each exposure measure. All specifications include year, state, and occupation fixed effects, demographic controls (sex, age, age^2, college, marital status, race), and person weights; standard errors are clustered at the occupation level.

3.1. Employment, wages, and hours

Tables 6–8 report DID estimates for employment, log wages, and usual weekly hours. Columns (1)–(3) use each continuous (z-scored) exposure measure separately. Column (4) is the preferred specification: it uses the mutually exclusive Automation/Augmentation dummies interacted with Post, with the Irrelevant group as the omitted baseline.

Table 6: DID: employed on AI exposure (Job Explorer, 2021–2024)

	(1) Total (z)	(2) Auto (z)	(3) Aug (z)	(4) Dummies
Total exposure × Post	0.000 (0.003)			
Automation exposure × Post		-0.001 (0.003)		
Augmentation exposure × Post			0.001 (0.002)	
Automation occ. × Post				0.003 (0.004)
Augmentation occ. × Post				-0.002 (0.002)
Observations	123,913	123,913	123,913	123,913
R-squared	0.043	0.043	0.043	0.043

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The employment margin (Table 6) displays no detectable response to ChatGPT exposure: the Automation and Augmentation dummies interacted with Post are both statistically indistinguishable from zero. This null result is consistent with the broader literature reviewed in Guarascio and Reljic (2025), which collects several papers finding that the early diffusion of Generative AI has not translated into visible reallocation of employment across occupations, as well as with the micro-level evidence in Humlum and Vestergaard (2024) and the macro projections in Cazzaniga et al. (2024). In the short run, layoffs are a costly and slow-moving adjustment margin, and workers in automation-prone occupations have had an incentive to adopt Generative AI *privately* to accelerate the very tasks that were at risk, thereby *preserving* their positions rather than losing them.

Table 7: DID: log_wage on AI exposure (Job Explorer, 2021–2024)

	(1) Total (z)	(2) Auto (z)	(3) Aug (z)	(4) Dummies
Total exposure × Post	-0.002 (0.009)			
Automation exposure × Post		0.015** (0.006)		
Augmentation exposure × Post			-0.006 (0.006)	
Automation occ. × Post				0.031*** (0.008)
Augmentation occ. × Post				-0.002 (0.004)
Observations	111,261	111,261	111,261	111,261
R-squared	0.168	0.168	0.168	0.168

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: DID: uhrswork on AI exposure (Job Explorer, 2021–2024)

	(1) Total (z)	(2) Auto (z)	(3) Aug (z)	(4) Dummies
Total exposure × Post	0.280 (0.209)			
Automation exposure × Post		0.422*** (0.036)		
Augmentation exposure × Post			0.115 (0.161)	
Automation occ. × Post				0.880*** (0.122)
Augmentation occ. × Post				0.529*** (0.038)
Observations	114,600	114,600	114,600	114,600
R-squared	0.076	0.076	0.076	0.076

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Wages (Table 7) show a clear asymmetry. Automation-dominant occupations experience a precisely estimated +3.1 log-point increase in wages after 2022 relative to the irrelevant group, while augmentation-dominant occupations show no significant change. A natural interpretation is that workers whose tasks are amenable to full delegation to a chatbot reaped the largest unilateral productivity gain from ChatGPT, in line with the productivity estimates of Brynjolfsson et al. (2023) and Czarnitzki et al. (2023). The absence of an augmentation effect is, I conjecture, partly a technology story: the GPT-3.5 model that was publicly released in November 2022 and most of GPT-4's 2023 iterations were better suited to one-shot delegation than to the iterative, feedback-loop interactions that characterise augmentation. More recent frontier models (for instance Claude 3.5/4, on which the AEI is built) have narrowed this gap, but the post-period largely pre-dates that turn.

Hours worked (Table 8) increase significantly in both exposure groups: Automation-dominant occupations work on average about 0.88 more hours per week after ChatGPT, and Augmentation-dominant occupations about 0.53 more, each relative to the irrelevant baseline. This is consistent with a view in which AI-exposed workers absorb the time saved on automatable tasks by expanding the scope of their work rather than reducing working time — an extensive-margin response also documented in survey evidence from Microsoft's *Work Trend Index* (Microsoft and LinkedIn, 2024) and discussed in the productivity literature (Brynjolfsson et al., 2023; Humlum and Vestergaard, 2024).

3.2. Heterogeneity

Tables 9–11 explore heterogeneity in the Automation/Augmentation dummy effects by gender, education, and race. The most consistent pattern is that *women* benefit disproportionately from the augmentation channel: the triple interaction Aug $\times$ Post $\times$ Female is positive and statistically significant for employment (+1.5 pp), log wages (+2.4 log points), and goes in the expected direction for hours. This gender tilt is consistent with the compositional evidence in Cazzaniga (2024), who document that women are over-represented in occupations where AI is complementary rather than substituting, and with the ILO exposure indices in Gmyrek et al. (2025b). By contrast, the college and white interactions are either insignificant or negative, suggesting that the early wage returns from augmentation are not concentrated among the most-educated or historically advantaged groups.

Table 9: Heterogeneity: employed

	(1) Female	(2) College	(3) White
Automation occ. $\times$ Post	0.004	0.002	-0.006
	(0.003)	(0.004)	(0.007)
Augmentation occ. $\times$ Post	-0.012***	0.004	0.004
	(0.003)	(0.003)	(0.003)
Auto $\times$ Post $\times$ female	-0.015		
	(0.015)		
Aug $\times$ Post $\times$ female	0.015***		
	(0.003)		
Auto $\times$ Post $\times$ college		0.002	
		(0.003)	
Aug $\times$ Post $\times$ college		-0.016***	
		(0.003)	
Auto $\times$ Post $\times$ white			0.014*
			(0.007)
Aug $\times$ Post $\times$ white			-0.008**
			(0.003)
Observations	123,913	123,913	123,913
R-squared	0.043	0.043	0.043

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Heterogeneity: log_wage

	(1) Female	(2) College	(3) White
Automation occ. × Post	0.031***	0.038***	0.014
	(0.008)	(0.008)	(0.024)
Augmentation occ. × Post	-0.019*	0.007	0.051**
	(0.009)	(0.008)	(0.020)
Auto × Post × female	0.001		
	(0.032)		
Aug × Post × female	0.024**		
	(0.008)		
Auto × Post × college		-0.048	
		(0.047)	
Aug × Post × college		-0.025	
		(0.026)	
Auto × Post × white			0.026
			(0.031)
Aug × Post × white			-0.081**
			(0.031)
Observations	111,261	111,261	111,261
R-squared	0.168	0.169	0.169

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 11: Heterogeneity: uhrswork

	(1) Female	(2) College	(3) White
Automation occ. × Post	0.899***	0.834***	0.618
	(0.080)	(0.112)	(0.433)
Augmentation occ. × Post	0.541***	0.651***	1.160***
	(0.097)	(0.073)	(0.188)
Auto × Post × female	-0.185		
	(0.583)		
Aug × Post × female	-0.017		
	(0.140)		
Auto × Post × college		0.342	
		(0.441)	
Aug × Post × college		-0.332	
		(0.203)	
Auto × Post × white			0.396
			(0.497)
Aug × Post × white			-0.959**
			(0.276)
Observations	114,600	114,600	114,600
R-squared	0.076	0.076	0.077

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

4. Conclusions

This paper studied the early labour-market impact of ChatGPT on the U.S. workforce by combining IPUMS ACS micro-data (2021–2024) with a revealed-preference measure of task-level AI exposure derived from Anthropic's Economic Index (Handa et al., 2025). For each of the 974 O*NET occupations covered by the Index I aggregated task-level automation and augmentation shares into a single occupation-level classification, producing three mutually exclusive groups — Automation-dominant, Augmentation-dominant, and Irrelevant — and used the last as the counterfactual in a difference-in-differences design built around the release of ChatGPT in November 2022.

The results paint a picture of early, heterogeneous, and mostly *positive* adjustments. Employment does not respond on the aggregate margin, consistent with the null findings in the survey-based literature collected by Guarascio and Reljic (2025) and Humlum and Vestergaard (2024); but wages in Automation-dominant occupations rise by about three log-points relative to the Irrelevant baseline, and usual weekly hours increase in both exposed groups. The heterogeneity analysis shows that women benefit disproportionately from the augmentation channel, with gains spanning wages, employment, and — directionally — hours. Taken together, these patterns suggest that the early adoption of Generative AI has, so far, been a productivity-enhancing technology for workers whose tasks *could* in principle be automated but who, in practice, seem to have used the tool privately to accelerate their own work rather than to be displaced by it. In the short run at least, the novel technology has not been a threat but an asset.

Three caveats qualify the contribution. First, the exposure measure is built from Claude conversations rather than ChatGPT conversations. The identification argument relies on the assumption that end-users do not meaningfully distinguish between Large Language Models at the task level, i.e. that the tasks Claude users delegate to AI in 2024–2025 are the same tasks ChatGPT users delegated to AI from late 2022 onwards. Direct ChatGPT usage data — such as the task-level panel documented in Chatterji et al. (2025) — would provide a cleaner mapping and is a natural next step for this line of work. Second, because the AEI is a post-treatment snapshot, the exposure assignment is effectively fixed at 2025 levels, which is restrictive if the task portfolio of LLMs shifted substantially between 2022 and 2025 —

plausibly toward more iterative, augmentation-style interactions, which may also explain the null wage effect in the augmentation group. Third, the specification relies on repeated cross-sections rather than a worker-level panel, so the mechanism discussion (private adoption and productivity offsets) is, at this stage, an interpretation of reduced-form moments rather than a direct test.

Despite these limitations, the combination of a revealed-preference exposure measure with a clean cut-off event offers a novel and empirically tractable way of studying the labour-market footprint of Generative AI — one that can be updated mechanically as new AEI releases, as well as future ChatGPT-specific data products, become available.

References

Acemoglu, D., Autor, D., 2011. Skills, tasks and technologies: Implications for employment and earnings, in: Handbook of Labor Economics. Elsevier. volume 4, pp. 1043–1171.

Acemoglu, D., Autor, D., Hazell, J., Restrepo, P., 2022. Artificial Intelligence and Jobs: Evidence from Online Vacancies. Journal of Labor Economics 40, S293–S340. doi:10.1086/718327.

Aldasoro, I., Gambacorta, L., Korinek, A., Shreeti, V., Stein, M., 2024. Intelligent financial system: How AI is transforming finance. BIS Working Paper .

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

Autor, D.H., Katz, L.F., Kearney, M.S., 2006. The Polarization of the U.S. Labor Market. American Economic Review 96, 189–194.

Autor, D.H., Katz, L.F., Krueger, A.B., 1998. Computing inequality: have computers changed the labor market? The Quarterly Journal of Economics 113, 1169–1214. URL: http://qje.oxfordjournals.org/.

Autor, D.H., Levy, F., Murnane, R.J., 2003. The skill content of recent technological change: An empirical exploration. The Quarterly Journal of Economics 118, 1279–1333.

Bick, A., Blandin, A., Deming, D.J., 2025. The rapid adoption of generative AI. National Bureau of Economic Research Working Paper .

Brynjolfsson, E., Li, D., Raymond, L.R., 2023. Generative AI at work. Quarterly Journal of Economics .

Cazzaniga, M., 2024. Gen-AI : Artificial Intelligence and the Future of Work. International Monetary Fund.

Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A.J., Pizzinelli, C., Rockall, E., Tavares, M.M., 2024. Gen-AI: Artificial intelligence and the future of work. IMF Staff Discussion Note .

Chatterji, A.K., Cunningham, T., Deming, D.J., Hitzig, Z., Ong, C., Shan, C.Y., Wadman, K., 2025. How people use ChatGPT: Task-level evidence from a representative panel of consumers. NBER Working Paper .

Crane, L.D., Soto, P., 2026. Generative AI and the labor market .

Czarnitzki, D., Fernández, G.P., Rammer, C., 2023. Artificial intelligence and firm-level productivity. Journal of Economic Behavior & Organization 211, 188–205.

Eloundou, T., Manning, S., Mishkin, P., Rock, D., 2024. GPTs are GPTs: Labor market impact potential of LLMs. Science 384, 1306–1308.

Felten, E.W., Raj, M., Seamans, R., 2023. How will language modelers like ChatGPT affect occupations and industries? arXiv preprint arXiv:2303.01157 .

Gmyrek, P., Berg, J., Bescond, D., 2025a. Generative AI and Jobs: An Analysis of Potential Effects on Global Employment. The Polish Journal of Economics 323, 6–30. doi:10.33119/gn/203716.

Gmyrek, P., Berg, J., Kamiński, K., Konopczyński, F., Ładna, A., Nafradi, B., Rosłaniec, K., Troszyński, M., 2025b. Generative AI and Jobs A Refined Global Index of Occupational Exposure. URL: https://doi.org/10.54394/HETP0387, doi:10.54394/HETP0387.

Guarascio, D., Reljic, J., 2025. Artificial intelligence, employment and wages: Evidence from European industries. Industrial and Corporate Change .

Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., Troy, K.K., Amodei, D., Kaplan, J., Clark, J., Ganguli, D., 2025. Which economic tasks are performed with AI? evidence from millions of claude conversations. Anthropic Economic Index Working Paper .

Haque, O.S., Li, X., 2025. Occupational heterogeneity in AI exposure .

Humlum, A., Vestergaard, E., 2024. The adoption of ChatGPT. University of Chicago Becker Friedman Institute Working Paper .

Microsoft and LinkedIn, 2024. Work trend index annual report: AI at work is here. now comes the hard part. Available at https://www.microsoft.com/en-us/worklab/work-trend-index/.

Zarifhonarvar, A., 2024. Economics of ChatGPT: A labor market view on the occupational impact of artificial intelligence. Journal of Electronic Business & Digital Economics .