

Lodefalk, Magnus; Löthman, Lydia; Koch, Michael; Engberg, Erik

Working Paper

Same storm, different boats: Generative AI and the age gradient in hiring

Working Paper, No. 2/2026

Provided in Cooperation with:

Örebro University School of Business

Suggested Citation: Lodefalk, Magnus; Löthman, Lydia; Koch, Michael; Engberg, Erik (2026) : Same storm, different boats: Generative AI and the age gradient in hiring, Working Paper, No. 2/2026, Örebro University School of Business, Örebro

This Version is available at:

<https://hdl.handle.net/10419/339320>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WORKING PAPER 2/2026 (ECONOMICS)

Same Storm, Different Boats: Generative AI and the Age Gradient in Hiring

Magnus Lodefalk, Lydia Löthman, Michael Koch and Erik Engberg

ISSN 1403-0586

Örebro University School of Business
SE-701 82 Örebro, Sweden

Same Storm, Different Boats: Generative AI and the Age Gradient in Hiring*

Magnus Lodefalk^{a,b,*}, Lydia Löthman^a, Michael Koch^{c,d}, Erik Engberg^{a,b}

^a*Örebro University School of Business, SE-701 82 Örebro, Sweden*

^b*RATIO Institute, Stockholm, Sweden*

^c*Department of Economics and Business Economics, Aarhus University, Denmark*

^d*Kiel Institut für Weltwirtschaft, Germany*

Abstract

We show that the age composition of employment within Swedish employers shifts after the arrival of generative AI, with no corresponding reduction in aggregate labour demand. Using 4.6 million job advertisements from Sweden's largest recruitment platform, we find that the broad decline in postings since 2022 aligns with monetary tightening rather than AI, exploiting Sweden's seven-month gap between the Riksbank's first rate hike and the launch of ChatGPT as a timing test. We then use full-population employer–employee register data and an employer-level difference-in-differences design to estimate how AI exposure affects employment composition across six age groups. An event study documents an accelerating decline in employment of 22–25-year-olds in high-AI-exposure occupations, reaching 5.5 per cent by early 2025 relative to less exposed occupations within the same employers, while employment of workers over 50 rose by 1.3 per cent. The widening age gradient suggests that generative AI reshapes hiring composition rather than aggregate demand, with the adjustment burden falling disproportionately on entry-level workers.

Keywords: Generative artificial intelligence, Job postings, Labour demand, Employment composition, Monetary policy

JEL: J23, J24, O33

*This version: 16 March 2026. Latest one at <https://www.ai-econlab.com/papers>.

*Corresponding author

Email address: magnus.lodefalk@oru.se (Magnus Lodefalk)

1. Introduction

A widely circulated chart shows US stock prices rising more than 70% since November 2022 while job openings fell more than 30%, suggesting “two economies” with AI as the suspected driver (Thompson, 2025). Using ADP payroll data, Brynjolfsson et al. (2025) document a 16% employment decline among young US workers in AI-exposed occupations, the “canaries in the coal mine”; Hosseini Maasoum and Lichtinger (2025) find a similar seniority gradient in US resume and job posting data.

We examine both the posting and the employment margin in Sweden and exploit a natural timing test: the Riksbank began raising rates in April 2022, seven months before ChatGPT launched in November 2022. If AI drives the differential posting decline, it should emerge only after late 2022; if monetary tightening is the primary driver, it should begin with the rate hike. Using 4.6 million job ads from Platsbanken (2020–2026), matched to the Dynamic AI Occupational Exposure (DAIOE) index (Engberg et al., 2024), we decompose the aggregate posting decline by AI exposure. The posting decline aligns with the Riksbank’s rate hike rather than AI, and concentrates in non-teleworkable occupations.

We complement the posting evidence with monthly population-register employment data, estimating employer-level difference-in-differences across six age groups following Brynjolfsson et al. (2025). Our results reveal that the employment gap for 22–25 year olds in high-AI-exposure occupations widens progressively after the ChatGPT launch, reaching 5.5 per cent by early 2025 relative to less exposed occupations.

We make three contributions. First, by examining both the posting and the employment margin, we reconcile the aggregate posting decline with a compositional employment shift: the broad posting decline reflects macroeconomic tightening, while AI appears to reshape hiring composition within employers. Second, Sweden’s seven-month gap between the Riksbank’s rate hike and ChatGPT provides a timing identification that separates monetary and AI channels, a timing decomposition not employed in existing designs, which address interest rates through cross-sectional sample splits (Brynjolfsson et al., 2025) or not at all (Kauhanen and Rouvinen, 2025). Third, we provide the first population-register evidence of an accelerating age-gradient “canaries” effect outside the US, contrasting with the Finnish null, which persists even under the same employer-level design (Kauhanen and Rouvinen, 2025, 2026).

The paper is organized as follows. Section 2 presents the data and empirical strategy, Section 3 reports the main results, and Section 4 concludes. We complement the main text with an Online Appendix, which documents the AI-exposure measure and its occupational distribution; provides robustness checks for the posting analysis, including macro-level contextualisation of the “scary chart”; presents the employer-level employment analysis in detail and compares our design and results to the emerging literature; and documents all data sources, coverage, and access procedures.

2. Data and empirical strategy

Job-posting margin. We use the full population of job advertisements published on Platsbanken from January 2020 through February 2026.¹ Each ad contains a publication date, SSYK 2012 four-digit occupation code, number of vacancies, municipality, and employer. After removing ads without valid SSYK codes and deduplicating on ad identifiers, we retain 4.6 million ads across 400 occupations. We aggregate these to occupation-by-month cells (26,672 observations).²

We measure generative AI (GAI) exposure using the DAIOE index (Engberg et al., 2024), which maps AI benchmark capabilities to Swedish occupations at the SSYK 4-digit level. We use the 2023 cross-section of the percentile ranking for generative AI and classify occupations into quartiles, where Q4 is the most exposed.

Our baseline specification is:

$$\ln(\text{postings}_{it}) = \alpha_i + \gamma_t + \beta_1 \cdot \text{PostRB}_t \cdot \text{High}_i + \beta_2 \cdot \text{PostGPT}_t \cdot \text{High}_i + \varepsilon_{it}, \quad (1)$$

where i indexes SSYK4 occupations, t indexes months, and α_i and γ_t are occupation and time fixed effects. $\text{PostRB}_t = 1$ if $t \geq$ April 2022 (the Riksbank’s first rate hike); $\text{PostGPT}_t = 1$ if $t \geq$ December 2022 (the first full month after ChatGPT’s November 30 launch); and $\text{High}_i = 1$ if occupation i is in the top quartile of genAI exposure. Standard errors are clustered at

¹Historical bulk data (2020–2025) from <https://data.jobtechdev.se/annonser/historiska/>; 2026 data from the real-time JobStream API.

²Platsbanken is the largest recruitment site in Sweden; until 2007, private and local-government employers were legally required to post vacancies there (Cronert, 2022). Any gap from platform migration biases us toward finding a negative AI effect on postings, making the null $\hat{\beta}_2$ conservative (see Online Appendix).

the occupation level. The key test is whether the relative posting decline in high-exposure occupations begins with monetary tightening ($\hat{\beta}_1$) or with the advent of generative AI ($\hat{\beta}_2$).

Employment margin. For employment, we use monthly employer declaration (AGI) register data from Statistics Sweden (SCB) covering the full population of employed individuals with SSYK 4-digit occupation, age, and gender.³ AI exposure is assigned from the DAIOE quartile of the worker’s most recent occupation.

We re-estimate Equation (1), using employment counts rather than postings, but now at the detailed employer \times AI-exposure-quartile \times month level and separately for each age group. Concretely, we replace the occupation fixed effects α_i with employer \times exposure-quartile fixed effects, and the month fixed effects γ_t with employer \times month fixed effects, following Brynjolfsson et al. (2025). These adjustments absorb all time-varying employer-level shocks, so identification comes solely from within-employer recomposition of employment across AI-exposure quartiles. The resulting coefficients – now called γ_2 – are age-specific analogues of β_2 in Equation (1), interpreted as the differential employment change in high-AI-exposure occupations within the same employer after the arrival of generative AI. The full specification is provided in the Online Appendix.

3. Results

Figure 1 shows the Swedish version of the “scary chart.”⁴ OMXS30 rose to over 175% of its February 2020 level by early 2026, while job postings across all AI exposure quartiles declined following the Riksbank rate hike in April 2022. The decline is broad-based across exposure quartiles, though Q4 declines somewhat more. The divergence between stock prices and postings opens seven months before ChatGPT, coinciding with the start of monetary tightening. The pattern is identical when using the OMXSPI All-Share index instead of the OMXS30 (see Online Appendix).

³The AGI system has been stable since 2019, and with no SSYK 2012 changes. Growth in platform-based self-employment (F-skatt) among young workers may lower measured employment but should not differentially affect AI-exposed occupations.

⁴Daily OMXS30 prices from Yahoo Finance are resampled to monthly averages and indexed to 100 at February 2020.

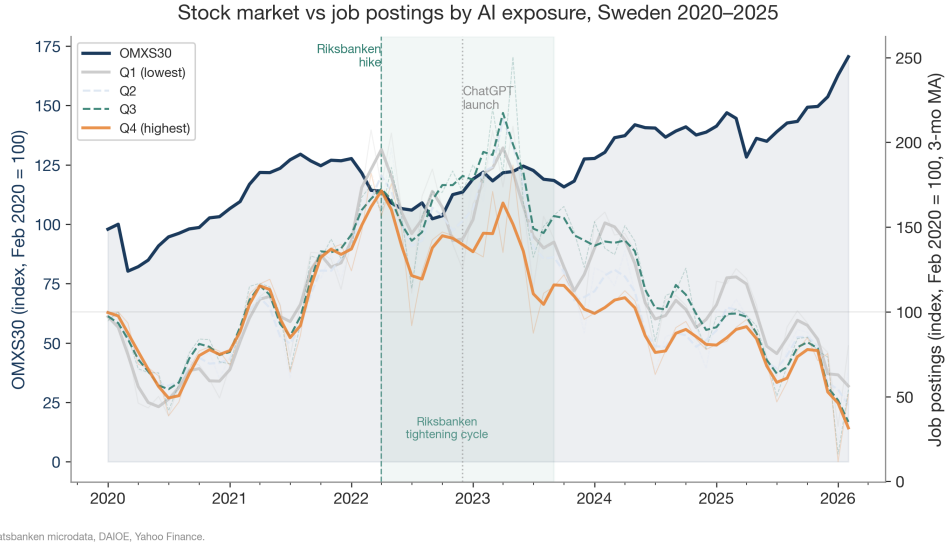


Figure 1: Stock market vs job postings by AI exposure quartile, Sweden 2020–2026. OMXS30 index (left axis) and Platsbanken posting index by genAI exposure quartile (right axis), both indexed to 100 at February 2020. Vertical lines mark the Riksbank’s first rate hike (April 2022) and ChatGPT launch (November 2022).

The difference-in-differences estimation of Equation (1) yields at most a small additional posting decline in high-AI-exposure occupations after ChatGPT ($\hat{\beta}_2 = -0.062$, $p = 0.11$), while the Riksbank rate-hike interaction is precisely estimated ($\hat{\beta}_1 = -0.127$, $p < 0.01$). AI exposure is uncorrelated with revealed monetary policy sensitivity ($r = 0.04$, $p = 0.51$; Online Appendix), confirming distinct channels. The coefficient is negative in seven of eight robustness specifications but statistically insignificant in six; results are directionally similar using the Eloundou et al. (2024) GPT exposure score, though the timing attribution shifts toward the ChatGPT period with that measure (Online Appendix). A teleworkability split (Dingel and Neiman, 2020) reveals that the ChatGPT coefficient is zero in teleworkable occupations ($\hat{\beta}_2 = -0.005$, $p = 0.92$) but significant in non-teleworkable ones ($\hat{\beta}_2 = -0.233$, $p < 0.01$), suggesting AI posting displacement concentrates where remote work cannot cushion the impact.

The employment margin analysis reveals nuanced differences across age groups. Figure 2 shows the pattern (3-month moving average; the unsmoothed series appears in the Online Appendix). Young workers (22–25)

in high-AI-exposure occupations diverge sharply from older workers after the ChatGPT launch, widening through 2025, while older workers (26+) are stable regardless of exposure.

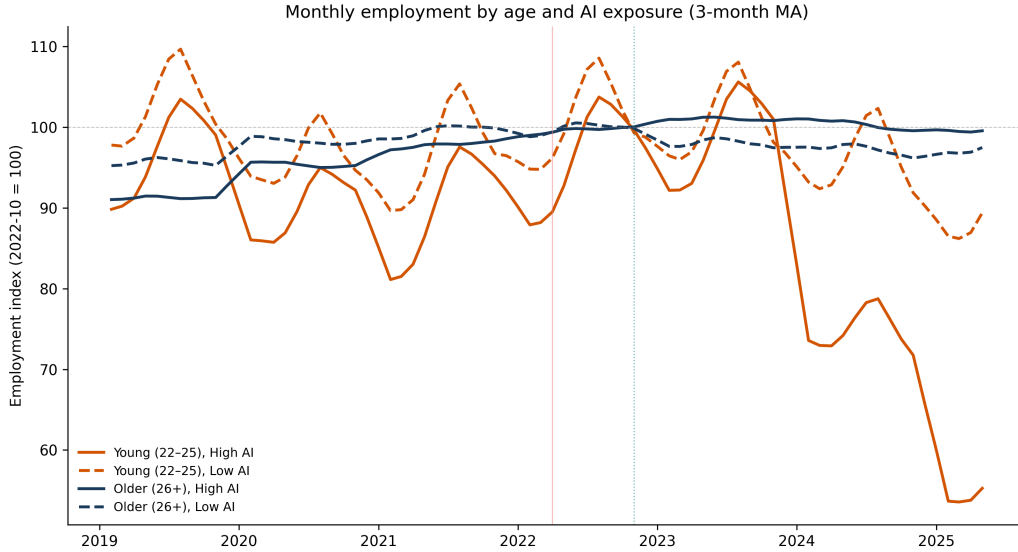


Figure 2: Monthly employment by age group and AI exposure, Sweden 2019–2025 (3-month moving average). AGI register data indexed to October 2022 = 100. Young (22–25) vs older (26+) workers, split by AI exposure quartile (DAIOE). Vertical lines mark the Riksbank’s first rate hike (April 2022) and ChatGPT launch (November 2022).

Figure 3 presents the event study. The average post-ChatGPT coefficient for 22–25 year olds is modest ($\hat{\gamma}_2 = -0.010$, $p < 0.01$), but the event study suggests a sharply accelerating divergence: by 2025H1, employment of young workers in top-quartile AI-exposed occupations has fallen 5.5 per cent relative to less exposed occupations, while employment of workers over 50 has risen 1.3 per cent.⁵ By 2025H1, the age gradient is monotonic: the 26–30 group

⁵Swedish bachelor’s graduates are at minimum 22, so the 22–25 bracket captures early-career workers. Pre-trend coefficients are jointly significant for all age groups, and pre-period magnitudes are comparable to the average post-ChatGPT coefficient, warranting caution. However, for the core 22–25 and 26–30 groups, the pre-trend slopes are positive (opposite to the post-treatment decline), so linear extrapolation works against the finding; the accelerating trajectory from 2024H1 onward is qualitatively distinct from the flat pre-period pattern; and a placebo test setting the treatment date to July 2022 (between the Riksbank hike and ChatGPT) shows flat coefficients through 2023 before the sharp

declines by 4.9 per cent, with progressively smaller effects for older groups (Online Appendix). The pattern is consistent with cumulative technological displacement diffusing up the age distribution rather than a single discrete shock. Only 13 per cent of young entrants first observed in 2023 are in top-quartile occupations, limiting the scope for a mechanical composition channel (Online Appendix). The effect is present for both men and women but approximately twice as large for young women ($\hat{\gamma}_2 = -0.016$ vs -0.007); reweighting women’s occupational distribution to match men’s narrows the gap by one-third, indicating that the remainder reflects within-occupation gender differences (Online Appendix).

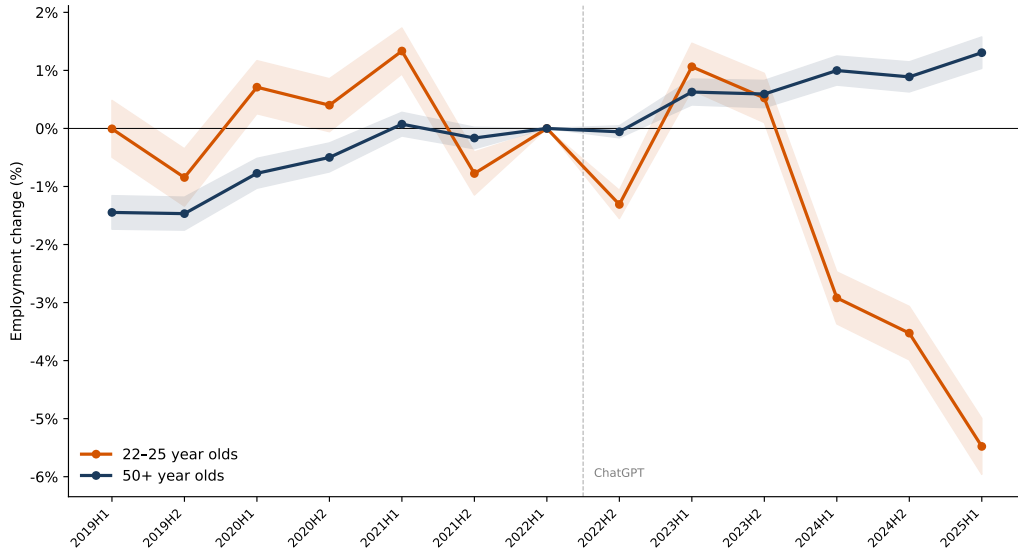


Figure 3: Event study of employment effects by age group and AI exposure, ages 22–25 and 50+. Coefficients from employer-level regressions with employer×quartile and employer×month fixed effects, half-year periods, reference period 2022H1. Shaded areas show 95% confidence intervals. Dashed line marks ChatGPT launch (November 2022). Coefficients from the log specification are interpreted as approximate percentage changes. All six age groups are shown in the Online Appendix.

acceleration in 2024–2025 (Online Appendix).

4. Conclusion – Implications, Limitations, and Policy Directions

Our analysis shows that the aggregate decline in job postings in Sweden aligns more closely with macroeconomic tightening than with generative AI exposure. Yet this aggregate pattern conceals a compositional shift that accelerates over time: an employer-level event study reveals a widening age gradient in the employment effects of AI exposure. By early 2025, young workers in AI-exposed occupations experienced large employment declines relative to less exposed occupations, while older workers saw modest gains. The finding is consistent with Brynjolfsson et al. (2025) in the US and the seniority-biased pattern in Hosseini Maasoum and Lichtinger (2025), contrasts with the Finnish null (Kauhanen and Rouvinen, 2025, 2026), which persists even under the same employer-level identification, and is complementary to the intensive-margin null in Humlum and Vestergaard (2025).⁶

Taken together, these results support the interpretation that generative AI reshapes the composition of employment rather than reducing aggregate labour demand. The observed age gradient maps onto the task-based framework (Acemoglu and Restrepo, 2019), in which AI substitutes for codified entry-level tasks while complementing the tacit judgement of experienced workers (Acemoglu et al., 2026): employers may post similar numbers of vacancies but fill them with fewer junior workers per position, consistent with a null on postings but a shift in employment composition.

Several caveats qualify the interpretation. First, pre-trend tests reject for all age groups, and pre-period magnitudes are comparable to the average post-ChatGPT coefficient, though for the core age groups the pre-trend direction is opposite to the post-treatment effect; a placebo date test confirms that the acceleration is specific to the post-ChatGPT period (Online Appendix). A formal sensitivity analysis following Rambachan and Roth (2023) would further quantify robustness to plausible trend violations. Second, we do not directly observe employers’ AI adoption, leaving room for unobservable organisational or strategic changes correlated with AI exposure. Third, population-level employment remains broadly stable, pointing to reallocation rather than destruction; our analysis does not track potential transitions to other occupations or sectors. Fourth, this remains an observational study:

⁶In the Online Appendix, we also provide an overview of recent studies examining how generative AI exposure affects employment or job postings differentially by worker age or seniority, and compare their findings with our own.

Swedish employment protection (LAS) creates seniority-based retrenchment incentives, anticipatory educational sorting by young workers, or differential career progression pathways could partly account for the age gradient, though the employer \times month fixed effects absorb employer-level cyclical effects.

The results imply that policymakers should focus on supporting early-career workers in AI-exposed occupations, where the adjustment burden appears greatest. Two priorities follow. First, educational and vocational systems should strengthen transition pathways into occupations where entry-level tasks are changing most rapidly. Second, employers may need guidance in redesigning junior roles to preserve opportunities for skill formation in an AI-augmented environment.⁷

AI disclosure

During the preparation of this work, the authors used Claude (Anthropic) for code development, data processing, and editorial suggestions. The authors reviewed and edited all content and take full responsibility for the publication.

Data availability

Platsbanken job advertisements are publicly available (CC0 licence). The employment analysis uses confidential administrative microdata from Statistics Sweden (SCB), accessed via the secure MONA platform; details on data access are in the Online Appendix. All replication code is available at <https://github.com/Magnus-L/canaries-sweden>.

Acknowledgements

We thank Lars Hultkrantz, Sune Karlsson, and participants at the AIS-CAF workshop (March 2026) for helpful comments. Lodefalk acknowledges financial support from the Torsten Söderberg Foundation (grants E46/21, ET3/23), and Lodefalk, Löthman and Engberg from WASP-HS (grant 805). Ethical approvals: 2021-05040, 2022-03330-02, 2024-01714-0, 2025-04205-02.

⁷The managerial implication – that employers should rethink junior roles rather than eliminate them – is discussed in “How Big a Threat Is AI to Entry-Level Jobs?” The Economist, 29 January 2026.

References

- Acemoglu, D., Autor, D., Johnson, S., 2026. Building Pro-Worker Artificial Intelligence. Working Paper 34854. National Bureau of Economic Research.
- Acemoglu, D., Restrepo, P., 2019. Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives* 33, 3–30. doi:10.1257/jep.33.2.3.
- Brynjolfsson, E., Chandar, B., Chen, R., 2025. Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Working Paper. Stanford Digital Economy Lab.
- Cronert, A., 2022. When the paper tiger bites: Evidence of compliance with unenforced regulation among employers in Sweden. *Regulation & Governance* 16, 1141–1159.
- 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., 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science* 384, 1306–1308. doi:10.1126/science.adj0998.
- Engberg, E., Görg, H., Lodefalk, M., Javed, F., Långkvist, M., Monteiro, N.P., Kyvik Nordås, H., Pulito, G., Schroeder, S., Tang, A., 2024. AI Unboxed and Jobs: A Novel Measure and Firm-Level Evidence from Three Countries. Discussion Paper 16717. IZA.
- Hosseini Maasoum, S.M., Lichtinger, G., 2025. Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Resume and Job Posting Data. Working Paper. Harvard University. SSRN 5425555.
- Humlum, A., Vestergaard, E., 2025. Large Language Models, Small Labor Market Effects. Working Paper 33777. National Bureau of Economic Research.
- Kauhanen, A., Rouvinen, P., 2025. Assessing early labour market effects of generative AI: Evidence from population data. *Applied Economics Letters* doi:10.1080/13504851.2025.2513973.

- Kauhanen, A., Rouvinen, P., 2026. AI Has Not Impacted the Youth Labor Market in Finland. Working Paper 135. ETLA Economic Research.
- Rambachan, A., Roth, J., 2023. A more credible approach to parallel trends. *Review of Economic Studies* 90, 2555–2591.
- Thompson, D., 2025. Is this the new ‘Scariest chart in the world’? Derek Thompson’s Substack. 23 October 2025. Accessed February 2026.

Online Appendix: Same Storm, Different Boats: Generative AI and the Age Gradient in Hiring

Lodefalk, Löthman, Koch, and Engberg

This version: 16 March 2026.

Latest version: <https://www.ai-econlab.com/papers>. Replication code: <https://github.com/Magnus-L/canaries-sweden>.

Contents

I	Data and measurement	2
II	Posting and macro analysis	4
III	Employment analysis	15
IV	Data documentation	40

This appendix provides supplementary analysis for the main paper. Part I documents the AI exposure measure and its occupational distribution. Part II presents robustness checks for the posting analysis, including macro contextualisation of the “scary chart.” Part III reports the employer-level employment analysis, including placebo date tests and extensive-margin diagnostics, and compares our approach with the emerging literature. Part IV documents data sources, coverage, and access procedures.

Part I: Data and measurement

I.1. Top and bottom occupations

Table A1 lists the ten most and ten least genAI-exposed occupations according to the DAIOE index.

Table A1: Most and least genAI-exposed occupations (DAIOE)

SSYK	Occupation	GenAI pctI
<i>Most exposed (top 10)</i>		
2641	Authors and related writers	99.9
2122	Statisticians	99.7
2121	Mathematicians and actuaries	99.4
2415	Economists	99.2
2512	Software and systems developers	98.9
2145	Chemical engineers	98.7
2111	Physicists and astronomers	98.5
2414	Securities traders and fund managers	98.2
2513	Game and digital media developers	98.0
2112	Meteorologists	97.8
<i>Least exposed (bottom 10)</i>		
9120	Vehicle, window and related cleaners	2.2
2653	Dancers and choreographers	2.0
8350	Ships' deck crew and related workers	1.8
7113	Concrete workers	1.5
8341	Agricultural and forestry machinery operators	1.3
9310	Construction labourers	1.1
8342	Earth-moving machinery operators	0.8
8111	Miners and quarry workers	0.6
7121	Roofers	0.3
3421	Professional athletes	0.1

I.2. Summary statistics

Table A2 presents summary statistics for the posting sample by AI exposure quartile. Table A3 presents corresponding statistics for the employer-level employment sample.

Table A2: Summary statistics: posting sample (occupation×month panel, 2019:10–2026:02).

	Q1	Q2	Q3	Q4	All
Occupations (SSYK4)	103	99	76	91	369
Occupation×month obs	7,239	7,234	5,524	6,675	26,672
Ads per month (mean)	99.5	197.6	199.7	181.3	167.3
Ads per month (SD)	213.4	566.2	525.7	367.6	438.3
Ads per month (median)	22	53	49	63	43
Vacancies per month (mean)	256.6	428.7	308.3	258.9	314.6
Vacancies per month (SD)	687.3	1,230.2	812.8	714.9	899.2
DAIOE genAI pctl (mean)	12.4	37.4	61.9	88.2	48.4

Q1 = lowest genAI exposure; Q4 = highest. Quartiles based on DAIOE percentile ranking.

Table A3: Summary statistics: employer-level employment sample (employer×quartile×month panel, 2019:01–2025:06).

	Value
Employers	510,213
Months	78
Non-zero panel cells	75,180,417
Min employer size	5
<i>Non-zero cells by age group</i>	
22–25	6,390,184
26–30	9,540,844
31–34	9,225,915
35–40	11,776,499
41–49	15,820,838
50+	22,426,137
<i>Person-months by quartile</i>	
Q1 (lowest exposure)	79,444,673
Q2	118,145,854
Q3	73,105,626
Q4 (highest exposure)	107,021,074

Balanced panel with zero-filled cells. Employers with ≥ 5 workers.

Restricted to employers observed in both Q4 and at least one of Q1–Q3.

Table A4: Panel structure by age group: balanced employer×quartile×month panel.

	22–25	26–30	31–34	35–40	41–49	50+
Employers	48,211	69,145	64,981	78,086	97,848	131,294
Emp×Q cells	153,467	211,520	189,405	224,501	279,628	383,279
Total panel cells	11.97M	16.50M	14.77M	17.51M	21.81M	29.90M
Zero-cell share	0.704	0.661	0.655	0.625	0.588	0.560
Mean cell size	1.77	2.20	2.16	2.60	3.17	4.29
Mean (if > 0)	5.97	6.49	6.25	6.93	7.68	9.75

Zero-cell share: fraction of employer×quartile×month cells with zero employment.

The high zero-cell share for young workers reflects smaller cohort sizes, not data quality.

Part II: Posting and macro analysis

II.1. Sweden vs United States comparison

Figure A1 shows side-by-side comparisons of the stock market versus job postings divergence in the United States and Sweden. The US panel uses the S&P 500 and Indeed Hiring Lab aggregate job postings index; the Swedish panel uses OMXS30 and Platsbanken microdata. Both countries exhibit a similar qualitative pattern of rising stock markets and falling postings since mid-2022.

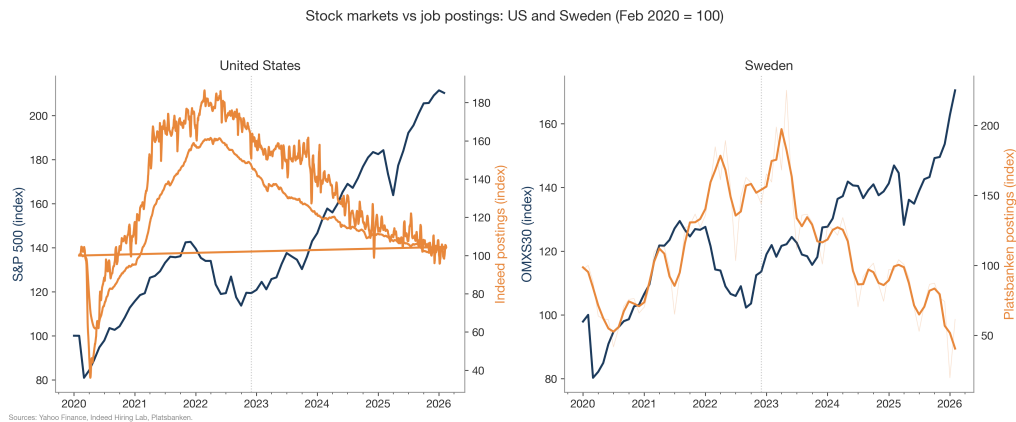


Figure A1: Stock markets vs job postings: United States (left) and Sweden (right), both indexed to 100 at February 2020.

II.2. Alternative stock market index (OMXSPI)

Figure A2 replicates the “scary chart” using the OMXSPI (OMX Stockholm All-Share Price Index), which covers all companies listed on Nasdaq Stockholm rather than only the 30 largest. The OMXS30 is dominated by banks, industrials, and a few technology firms, raising the concern that the stock market–postings divergence reflects the performance of a narrow set of large caps. The OMXSPI comparison shows that the pattern is virtually identical, confirming that the divergence is not an artefact of index composition.

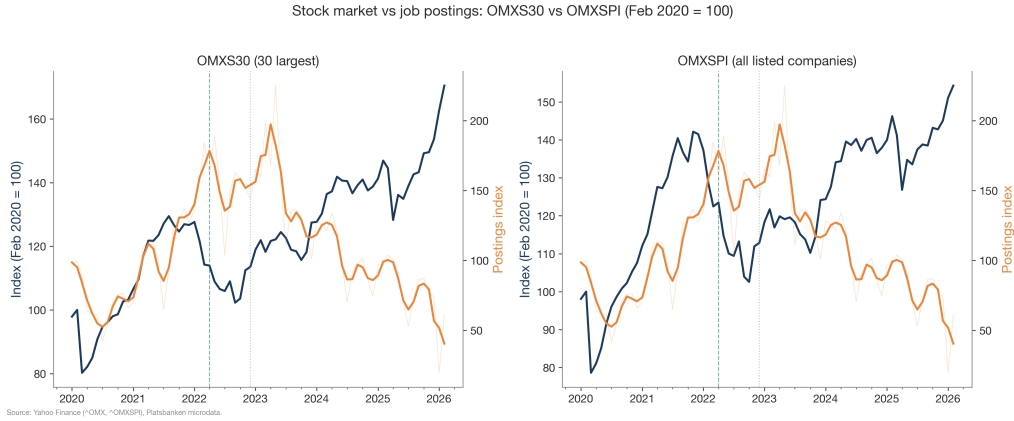


Figure A2: Stock market vs job postings using OMXS30 (left) and OMXSPI All-Share (right). Both panels show the same divergence pattern, confirming it is not driven by the composition of the OMXS30.

II.3. Individual quartile trends

Figure A3 shows each AI exposure quartile’s posting trajectory individually. All four quartiles exhibit similar cyclical patterns, peaking in mid-2022 and declining through 2023–2024. The parallelism of these trends is consistent with a common macroeconomic driver rather than AI-specific displacement.

II.4. Posting DiD and alternative specifications

Table A5 reports the main difference-in-differences estimates discussed in the text (Equation 1 of the main paper).

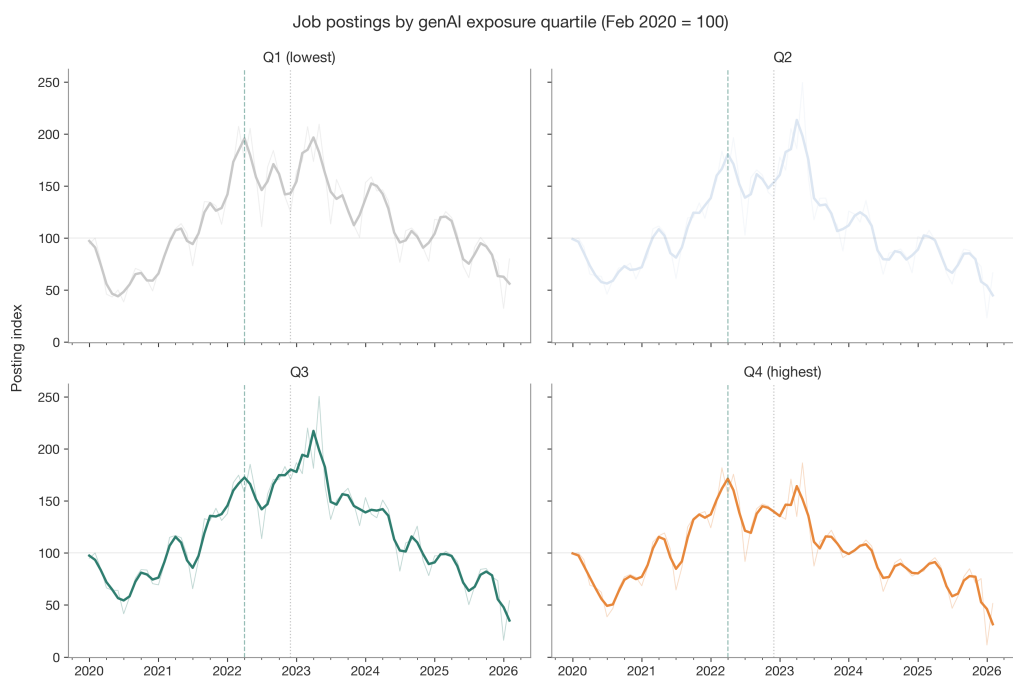


Figure A3: Job posting index by genAI exposure quartile (Feb 2020 = 100), individual panels.

Table A5: Difference-in-differences: postings by genAI exposure

	(1)	(2)	(3)	(4)
	Monetary	+ ChatGPT	+ Trends	Group×Time
Post-Riksbank × High	-0.178*** (0.041)	-0.127*** (0.039)	-0.068 (0.050)	-0.039 (0.049)
Post-ChatGPT × High		-0.062 (0.038)	0.018 (0.040)	-0.032 (0.048)
Time trend × High			-0.003** (0.001)	
Occupation FE	Yes	Yes	Yes	Yes
Month FE	Yes	Yes	Yes	
Occ. group × month FE				Yes
Dep. var. mean			3.77	
Occupations			369	
Observations			26,672	

Notes: Dependent variable: $\ln(\text{postings}_{it})$. High = top quartile of DAIOE genAI percentile ranking. Post-Riksbank = April 2022 onward. Post-ChatGPT = December 2022 onward. Column (4) replaces common month FE with SSYK 1-digit occupation group × month FE. Standard errors clustered at occupation level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table A6 reports results from alternative specifications: (i) using the all-apps AI exposure measure instead of genAI; (ii) using vacancy-weighted postings; (iii) excluding pandemic months (January–June 2020); (iv) using tercile classification instead of quartiles; (v) excluding IT/tech occupations (SSYK 25xx), following Brynjolfsson et al. (2025); (vi) restricting to a balanced panel of 308 occupations observed in every month; (vii) using language-modelling task exposure from DAIOE; (viii) adding quadratic occupation-specific time trends. The ChatGPT coefficient $\hat{\beta}_2$ is insignificant in six of eight specifications but significant at the 5% level in the all-apps measure ($\hat{\beta}_2 = -0.091$, $p = 0.018$) and tercile classification ($\hat{\beta}_2 = -0.081$, $p = 0.026$). The coefficient is negative in seven of eight specifications. The quadratic trend specification (viii) confirms that non-linear differential dynamics do not drive the results: the quadratic term is insignificant ($\hat{\delta}_2 = -0.00002$, $p = 0.56$) and the ChatGPT coefficient remains insignificant ($\hat{\beta}_2 = 0.019$, $p = 0.63$).

Software replication. The posting DiD is estimated in Python (`statsmodels`) and independently replicated in Stata 18.5 using `reghdfe`; point estimates match to three decimal places across all specifications. As a further check, a co-author independently replicated the main posting result using a separately constructed dataset of Platsbanken advertisements, generated with different code and downloaded at a different time. The two datasets correlate at $R^2 = 0.998$ at the occupation–month level, and the regression coefficients are robust across all approaches. For the employer-level employment DiD, the primary `PanelOLS` estimates (Approach A: absorbed fixed effects) are verified against a manual within-transformation (Approach B: sequential demeaning by employer \times quartile and employer \times month), with point estimates matching. Standard errors differ slightly between approaches due to different clustering levels (employer \times quartile in Approach A; employer in Approach B).

II.5. Alternative AI exposure measure: Eloundou et al. (2024)

To test sensitivity to the choice of AI exposure measure, we re-estimate specifications (2)–(4) using the GPT exposure score (β -measure) from Eloundou et al. (2024) instead of DAIOE. The β -measure combines human expert and GPT-4 ratings of task-level exposure to large language models ($E_1 + 0.5 \times E_2$, where E_1 is direct exposure and E_2 is exposure via complementary tools). We crosswalk from O*NET-SOC to SSYK 2012 via ISCO-08

Table A6: Robustness checks

Specification	$\hat{\beta}_1$ (Riksbank)	$\hat{\beta}_2$ (ChatGPT)	N	Occ.
Baseline (genAI Q4)	-0.1271*** (0.0388)	-0.0615 (0.0380)	26,672	369
All-apps measure	-0.0921** (0.0386)	-0.0911** (0.0385)	26,672	369
Vacancy-weighted	-0.1038** (0.0443)	-0.0612 (0.0437)	26,672	369
Excl. pandemic	-0.1159*** (0.0356)	-0.0617 (0.0380)	24,479	369
Terciles	-0.0936*** (0.0360)	-0.0808** (0.0362)	26,672	369
Excl. IT/tech	-0.1291*** (0.0406)	-0.0662* (0.0400)	26,148	362
Balanced panel	-0.1130*** (0.0377)	-0.0325 (0.0345)	22,176	308
Language model	-0.1154*** (0.0386)	-0.0564 (0.0387)	26,672	369
Quadratic trends	-0.0884*** (0.0323)	0.0193 (0.0402)	26,672	369

Notes: All specifications include occupation and month fixed effects. Standard errors (in parentheses) clustered at occupation level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

using the BLS and SCB classification bridges, averaging scores across many-to-many links. The crosswalk matches 341 of 369 occupations in our panel (92%).

Table A7 reports the results. The two measures agree on the direction: both yield negative Post-ChatGPT coefficients in the baseline specification. However, the timing attribution differs. DAIOE loads the effect onto the Riksbank period ($\hat{\beta}_1 = -0.127$, $p < 0.01$), while the Eloundou measure loads it onto the ChatGPT period ($\hat{\beta}_2 = -0.096$, $p = 0.025$). This likely reflects what each measure captures: DAIOE covers a broader range of AI capabilities in the area of language modelling and image generation, including pre-GPT applications, while the Eloundou β -measure is calibrated specifically to GPT-style generative AI. Both coefficients attenuate toward zero with occupation-specific trends or group \times month fixed effects, consistent with the fragility documented in the main analysis. Quartile assignments agree for 85% of occupations.

Table A7: Robustness: DAIOE vs. Eloundou et al. (2024) GPT exposure

	DAIOE			Eloundou β		
	(1)	(2)	(3)	(4)	(5)	(6)
Post-RB \times High	-0.127*** (0.039)	-0.068 (0.050)	-0.039 (0.049)	-0.022 (0.043)	0.029 (0.057)	0.037 (0.051)
Post-GPT \times High	-0.062 (0.038)	0.018 (0.040)	-0.032 (0.048)	-0.096** (0.043)	-0.028 (0.049)	-0.066 (0.049)
Occ. trends	Yes			Yes		
Occ. group \times month FE	Yes			Yes		
Occupations	369			341		
Observations	26,672			24,640		

Notes: Dependent variable: $\ln(\text{postings}_{it})$. All specifications include occupation and month FE (cols. 3 and 6 replace common month FE with occupation group \times month FE). “High” = top quartile. Post-RB = April 2022+; Post-GPT = December 2022+. Eloundou β : GPT exposure from Eloundou et al. (2024), crosswalked via SOC→ISCO→SSYK. Clustered SEs. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

II.6. Event study (monthly)

Figure A4 plots monthly DiD coefficients from an event-study specification, interacting month dummies with the high-exposure indicator (omitting February 2020 as the reference period). The pre-period coefficients fluctuate around zero with no systematic trend, though a joint Wald test rejects the

null that all 26 pre-Riksbank coefficients are zero ($\chi^2_{26} = 107.1, p < 0.01$). Aggregating to quarterly frequency reduces the number of pre-period coefficients to 9 but the test still rejects ($\chi^2_8 = 52.8, p < 0.01$), indicating genuine differential macro sensitivity across AI exposure groups rather than monthly noise. As Roth (2022) shows, a failing pre-test does not invalidate the design but does counsel caution; it motivates specification (4) in the main table, which conditions on occupation group \times month fixed effects, and the employer-level analysis, which provides a fundamentally stronger identification strategy. The post-ChatGPT coefficients show no additional structural break beyond the Riksbank hiking cycle in either specification.

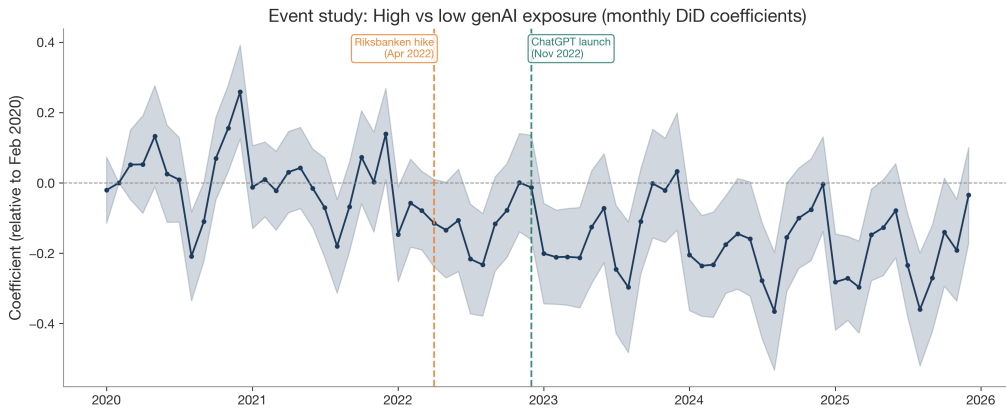


Figure A4: Event study: monthly DiD coefficients for high vs low genAI exposure occupations (reference: February 2020). Shaded area shows 95% confidence interval. Dashed lines mark the Riksbank’s first rate hike (April 2022) and ChatGPT launch (November 2022).

II.7. Interest rate sensitivity and AI exposure

A key concern for identification is whether AI-exposed occupations are also more sensitive to monetary policy, which would confound the rate-hike and ChatGPT interactions. To test this, we compute a “revealed” monetary policy sensitivity measure: the log change in mean postings between the pre-hike window (January–March 2022) and the post-hike, pre-ChatGPT window (May–November 2022). This window is uncontaminated by AI effects (ChatGPT launched November 30, 2022), so any posting decline is macroeconomic by definition.

Figure A5 plots this revealed rate sensitivity against DAIOE genAI exposure. There is no significant correlation (unweighted $r = 0.035$, $p = 0.507$; posting-weighted $r = -0.020$), indicating that the rate hike affected postings broadly across all AI exposure levels. This mirrors Brynjolfsson et al. (2025), who found a negative correlation between AI exposure and interest rate sensitivity in the US using the Zens et al. (2020) measure; Sweden shows essentially zero correlation with our revealed-preference approach.

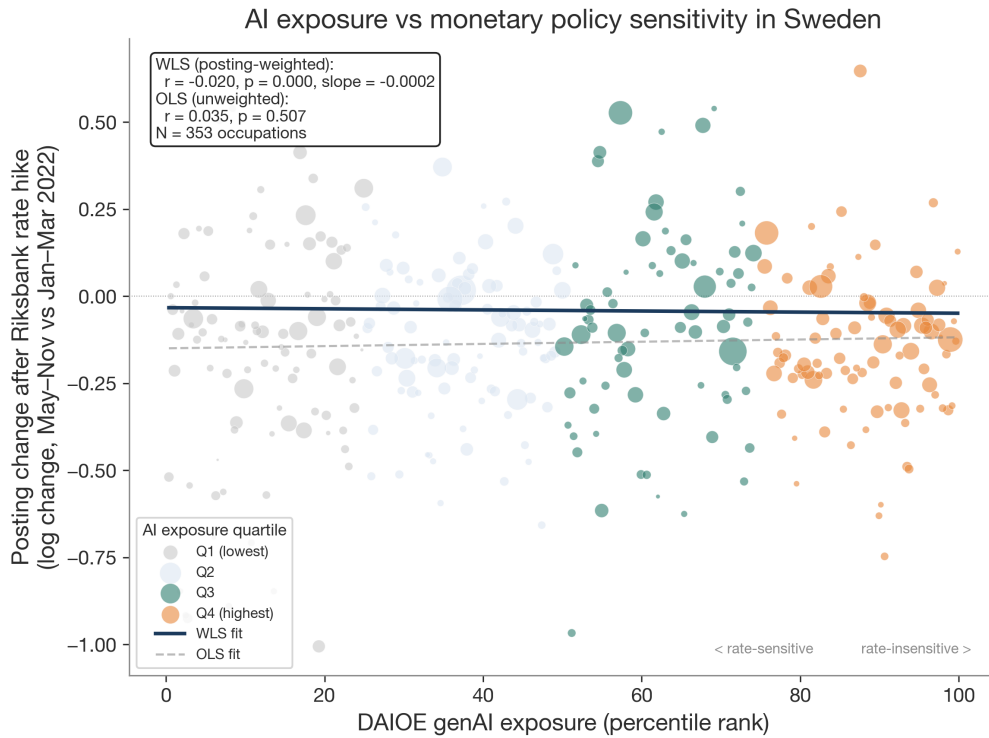


Figure A5: AI exposure (DAIOE genAI percentile rank) vs revealed monetary policy sensitivity (log change in postings, May–November vs January–March 2022). Bubble size proportional to total posting volume. No significant correlation: the rate hike affected all AI exposure levels similarly.

II.8. Teleworkability robustness (Dingel-Neiman)

AI-exposed occupations overlap substantially with teleworkable occupations (Dingel and Neiman, 2020). To test whether the posting null reflects

teleworkability rather than AI, we crosswalk the [Dingel and Neiman \(2020\)](#) teleworkability classification from SOC 2010 to SSYK 2012 via ISCO-08 and re-run specification (2) separately for teleworkable and non-teleworkable occupations (split at the median score).

Figure A6 reveals heterogeneity. In teleworkable occupations (187 SSYK codes), $\hat{\beta}_2$ is essentially zero (-0.005 , $p = 0.917$); the posting decline is entirely captured by the Riksbank interaction ($\hat{\beta}_1 = -0.138$, $p < 0.01$). In non-teleworkable occupations (182 SSYK codes), $\hat{\beta}_2 = -0.233$ ($p < 0.01$), indicating a significant additional decline after ChatGPT in high-AI-exposure occupations that cannot adapt through remote work. This suggests that AI displacement of postings, where it occurs, concentrates in occupations where remote work arrangements cannot cushion the impact.

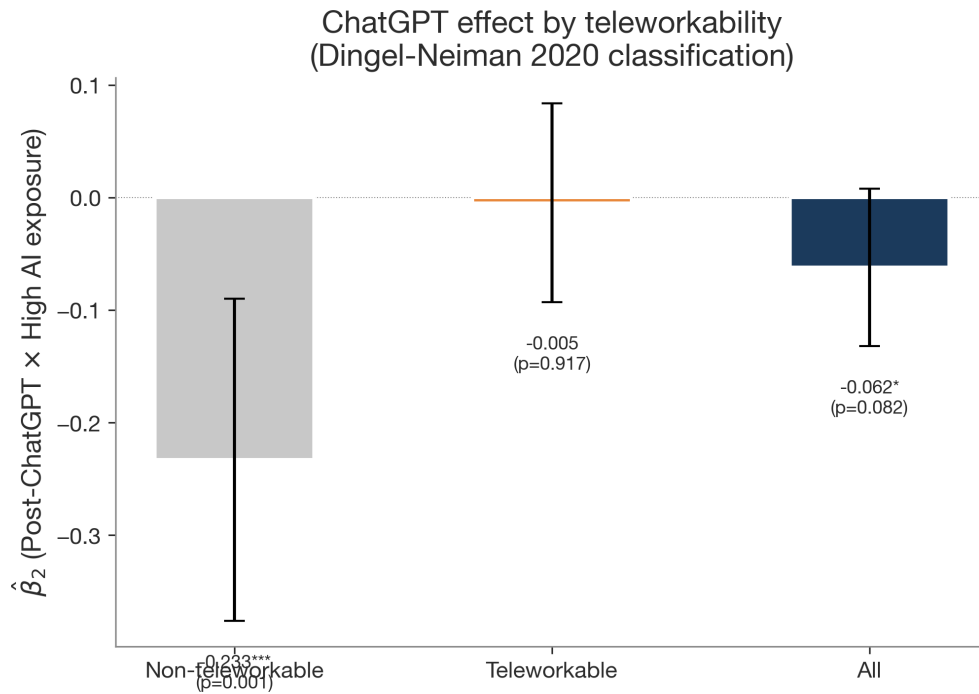


Figure A6: ChatGPT coefficient ($\hat{\beta}_2$) from specification (2) estimated separately for teleworkable and non-teleworkable occupations (Dingel-Neiman 2020 classification, crosswalked to SSYK 2012 via ISCO-08). Error bars show 95% confidence intervals. The AI posting effect concentrates in non-teleworkable occupations.

II.9. Sensitivity to violations of parallel trends

Figure A7 reports a sensitivity analysis following the relative magnitudes framework of Rambachan and Roth (2023). The average post-ChatGPT event-study coefficient for high-exposure occupations is $\hat{\theta} = -0.169$ (SE = 0.059), significantly negative under exact parallel trends ($\bar{M} = 0$). However, the “breakdown value” is $\bar{M} = 0.25$: if post-period violations of parallel trends are as little as 25% of the maximum pre-period violation, the honest confidence interval includes zero. This confirms that while there is suggestive evidence of a negative AI effect on postings, the finding is fragile, consistent with the imprecision documented in the main regression table.

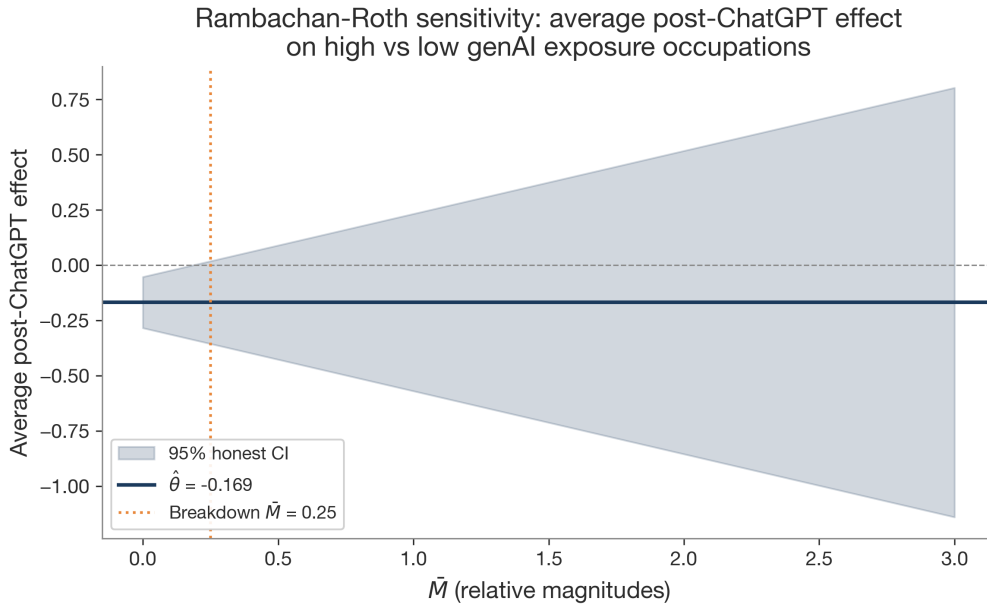
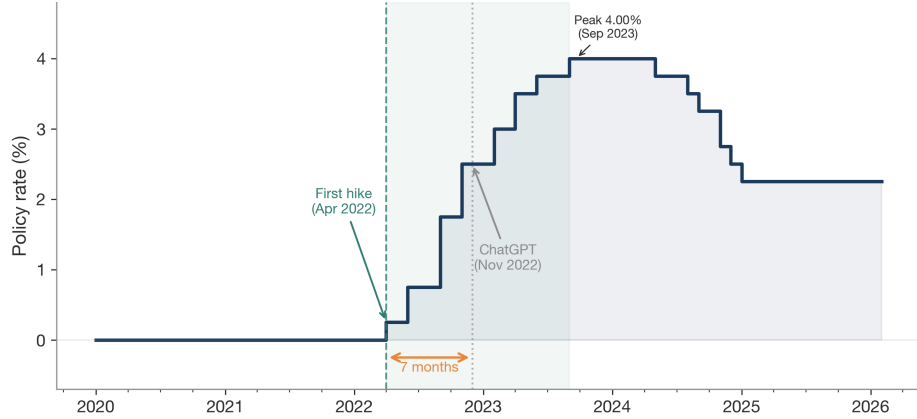


Figure A7: Rambachan-Roth sensitivity analysis for the average post-ChatGPT effect on high vs low genAI exposure occupations. The solid line shows the point estimate ($\hat{\theta} = -0.169$); the shaded area shows the 95% honest confidence interval as a function of \bar{M} (the maximum ratio of post- to pre-period trend violations). The dotted line marks the breakdown value $\bar{M} = 0.25$.

II.10. Riksbank policy rate timeline

Figure A8 shows the Riksbank’s policy rate from 2020 through early 2026. The tightening cycle ran from April 2022 (first hike, 0% to 0.25%) through

September 2023 (peak at 4.00%), followed by gradual easing beginning in May 2024. The 7-month gap between the first hike and ChatGPT’s launch (November 2022) is the identification window that separates monetary tightening from AI exposure in our design.



Source: Riksbank press releases (riksbank.se). Manually compiled and verified.

Figure A8: Riksbank policy rate, 2020–2026. The tightening cycle ran from April 2022 to September 2023 (0% to 4.00%), with easing from May 2024. The 7-month gap between the first hike and ChatGPT’s launch is marked.

Part III: Employment analysis

III.1. Employment by age group and AI exposure

We use monthly employer declarations (AGI) to the Swedish tax authority. These register data are available via SCB’s secure research environment (MONA). The AGI provides population-level individual employment records matched to SSK 4-digit occupations by age and gender (detailed documentation in Part IV). An employer-level difference-in-differences following Brynjolfsson et al. (2025) yields a modest average post-ChatGPT coefficient for the youngest workers ($\hat{\gamma}_2 = -0.010$, $p < 0.01$), but a half-year event study reveals an accelerating divergence: by 2025H1, employment of 22–25 year olds in top-quartile occupations has fallen by 5.5 per cent relative to less exposed occupations, while employment of workers over 50 has risen by 1.3 per cent. Swedish college graduation age is high by international standards (median 27–28), so the 22–25 bracket captures very-early-career workers.

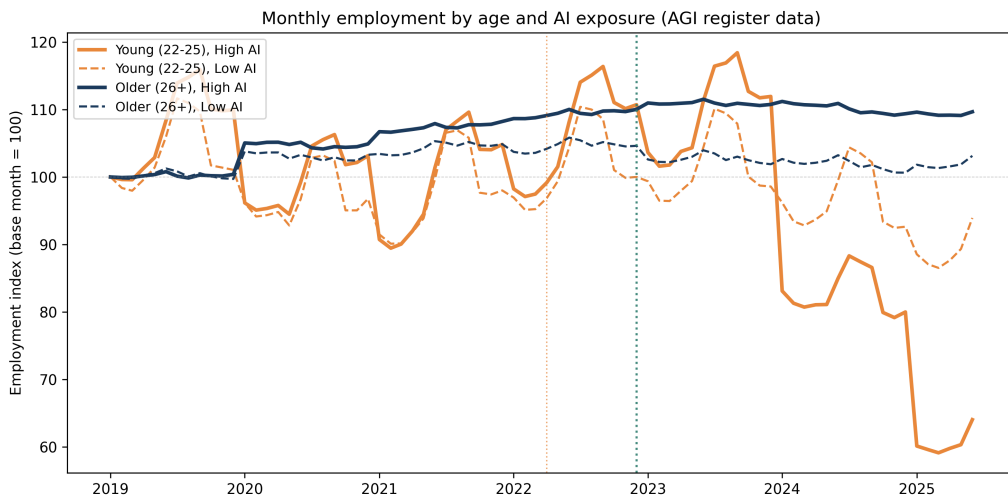


Figure A9: Monthly employment by age group and AI exposure, unsmoothed. Same data as Figure 2 in the main paper but without the 3-month moving average, showing that the divergence is not an artefact of smoothing. Strong seasonality is visible, particularly for the 22–25 group (summer employment peaks). The accelerating decline of young workers in high-AI occupations (solid orange) from 2024 onward partly reflects SSK matching attrition as LISA coverage ends in 2023 (see Part IV).

Reconciliation with the Finnish and Danish evidence. Our employer \times quartile and employer \times month fixed effects isolate within-employer recomposition across AI exposure quartiles. An earlier Finnish study finds null effects using occupation-level synthetic DiD (Kauhanen and Rouvinen, 2025). Importantly, Kauhanen and Rouvinen (2026) replicate the Brynjolfsson firm-level Poisson design with firm \times quintile and firm \times month fixed effects, estimated separately by age group, and still find no displacement among young Finnish workers – if anything, employment in more exposed quintiles *increased* slightly for early-career groups. The Finnish null therefore cannot be attributed to a weaker identification strategy: it persists under the same within-employer design that reveals a canaries effect in the US and Sweden. In Denmark, Humlum and Vestergaard (2025) find precise nulls on incumbents’ earnings and hours (the intensive margin), complementing our extensive-margin (composition) finding. Whether the Sweden–Finland divergence reflects differences in AI adoption intensity, labour market institutions, or industrial structure remains an open question.

Military conscription. Sweden reactivated military conscription (*värnplik*) in 2017, with annual intake expanding from approximately 5,000 in 2019 to 7,000 in 2024 and a target of 8,000 by 2025 (see [Plikt- och prövningsverket, 2025](#)). Roughly three-quarters of conscripts are male, and service lasts approximately 11 months at ages 18–19. Since military occupations carry no DAIOE score, expanding conscription mechanically reduces the civilian labour force participation of young men in our data, potentially confounding the canaries effect. Two observations mitigate this concern. First, the conscription expansion is gradual and affects both high- and low-AI-exposure occupations equally, since conscripts are drawn from across the occupational distribution. Second, our preferred specification following Brynjolfsson et al. (2025) restricts the young-worker sample to ages 22–25, above the conscription age, effectively eliminating this channel.

III.2. Employer-level difference-in-differences

The descriptive evidence in Figure 2 of the main paper and the analysis in the preceding section establish that young workers in high-AI-exposure occupations experienced disproportionate employment declines. However, this occupation-level analysis cannot rule out that unobserved employer-level

shocks – for instance, differential retrenchment by large employers concentrated in AI-exposed occupations – drive the pattern. To tighten identification, we estimate an employer-level difference-in-differences design closely following [Brynjolfsson et al. \(2025\)](#).

Estimating equation. For each age group $a \in \{22\text{--}25, 26\text{--}30, 31\text{--}34, 35\text{--}40, 41\text{--}49, 50+\}$, we estimate:

$$\ln(\text{emp}_{fqt}^a) = \alpha_{fq} + \beta_{ft} + \gamma_1 \cdot \text{PostRB}_t \cdot \text{HighQ4}_q + \gamma_2 \cdot \text{PostGPT}_t \cdot \text{HighQ4}_q + \varepsilon_{fqt} \quad (1)$$

where f indexes employers, $q \in \{1, 2, 3, 4\}$ indexes DAI OE genAI exposure quartiles, and t indexes year-months. α_{fq} are employer \times quartile fixed effects and β_{ft} are employer \times month fixed effects. The dependent variable is the log count of workers of age group a employed by employer f in occupations belonging to quartile q at time t . $\text{PostRB}_t = 1$ if $t \geq$ April 2022; $\text{PostGPT}_t = 1$ if $t \geq$ December 2022; $\text{HighQ4}_q = 1$ if the occupation belongs to the top quartile of genAI exposure. Standard errors are clustered at the employer \times quartile level.

Identification. The employer \times month fixed effects β_{ft} are the key identification device. They absorb *all* time-varying employer-level shocks – including employer-specific responses to monetary tightening, demand shocks, management changes, or strategic decisions – leaving only *within-employer, within-month* variation across AI exposure quartiles. The identifying assumption is that, absent AI, employers would have adjusted employment in high- and low-exposure occupations proportionally. This is substantially weaker than the parallel trends assumption required by the occupation-level analysis, which requires that aggregate shocks affect all occupations equally. Our quartile discretisation converts the continuous DAI OE measure into a binary treatment (Q4 vs Q1–Q3), avoiding the negative-weighting problems that [Callaway et al. \(2024\)](#) identify for continuous-dose TWFE designs; for a recent survey, see [Roth et al. \(2023\)](#).

Canaries test. The coefficient $\hat{\gamma}_2$ in Equation (1) measures the differential employment change in high-AI-exposure occupations (relative to lower quartiles, within the same employer, in the same month) after ChatGPT. By estimating the equation separately for each age group, we test whether this within-employer recomposition is age-graded. If $\hat{\gamma}_2$ is significantly negative for the youngest workers (22–25) but not for older groups, this constitutes

evidence of a canaries effect: employers differentially reduce young hires in AI-exposed roles after the advent of generative AI.

Event study. To validate the parallel trends assumption and trace the dynamics of any effect, we also estimate an event-study specification:

$$\ln(\text{emp}_{fqt}^a) = \alpha_{fq} + \beta_{ft} + \sum_{s \neq s_0} \delta_s \cdot \mathbf{1}[t \in s] \cdot \text{HighQ4}_q + \varepsilon_{fqt} \quad (2)$$

where s indexes half-year periods and $s_0 = 2022\text{H1}$ is the reference period (the half-year in which the Riksbank began its tightening cycle; since AI exposure is uncorrelated with monetary policy sensitivity, this partial overlap does not bias the pre-trend coefficients). As a robustness check, we also estimate the event study with $s_0 = 2021\text{H2}$ as the reference period, which avoids any overlap with the monetary tightening cycle; results are qualitatively unchanged. Pre-treatment coefficients $\hat{\delta}_s$ for $s < s_0$ should be statistically indistinguishable from zero if the parallel trends assumption holds within employers.

Sample construction. We use the same monthly AGI register data as the descriptive analysis, restricting to employers with at least five employee-month observations over the sample period and at least one worker observed in both high (Q4) and low (Q1–Q3) exposure occupations. This restriction ensures that identification comes from within-employer comparisons; single-quartile employers contribute no variation. Each employer-quartile-month cell records the count of workers in the relevant age group. The panel is balanced: cells with zero employment are included (via $\ln(0 + 1) = 0$) to capture extensive-margin adjustment, ensuring that employers shedding all workers of a given age group in high-exposure occupations contribute to identification rather than being dropped.

Estimation. The two-way high-dimensional fixed effects are absorbed using the method of alternating projections (`PanelOLS` in Python; `fixest` in R). With potentially millions of employer \times month fixed effects, memory constraints may bind; we implement a three-level estimation fallback: (i) direct absorption via `linearmodels/fixest`, (ii) manual within-transformation (de-meaning), (iii) occupation-level aggregation if employer-level estimation is infeasible. [Brynjolfsson et al. \(2025\)](#) use Poisson QMLE following [Chen and Roth \(2024\)](#); our OLS specification on $\ln(\text{employment} + 1)$ yields similar estimates when cell counts are large, as in our data. We plan to report Poisson

QMLE estimates as a robustness check in a revision. The +1 transformation retains employer×quartile×month cells with zero employment, which is informative for the extensive margin (whether an employer retains *any* workers of a given age group in high-exposure occupations). The posting specification uses $\ln(\text{postings})$ without the +1 adjustment, since zero-posting cells represent genuinely absent demand and the log transformation naturally drops them. Unlike the ADP payroll data in Brynjolfsson et al. (2025), which retains 3.5 to 5 million US workers per month after excluding part-time workers and those without job titles, our AGI register data cover the full population of Swedish employees, eliminating sample selection concerns.

Results. Table A8 reports the employer-level DiD estimates for each age group. The average post-ChatGPT coefficient $\hat{\gamma}_2$ is significantly negative only for the youngest workers (-0.010 for ages 22–25, $p < 0.01$), null for ages 26–30 ($+0.000$, $p = 0.90$), and positive for workers over 50 ($+0.010$, $p < 0.01$). The modest average for 22–25 year olds reflects the accelerating trajectory visible in the event study (below): the effect is near zero in the first year after ChatGPT and widens sharply from 2024H1 onward. The Riksbank coefficient $\hat{\gamma}_1$ is significantly negative for the youngest group (-0.012 , $p < 0.01$), indicating some differential monetary policy sensitivity for young workers in AI-exposed occupations. Since the employer×month fixed effects already absorb aggregate monetary effects, $\hat{\gamma}_1$ captures only the *differential* sensitivity of high-AI occupations to tightening. The presence of a significant $\hat{\gamma}_1$ makes $\hat{\gamma}_2$ conservative: it measures the additional AI-related shift beyond any monetary vulnerability channel. Moreover, the event study shows a qualitative acceleration from 2024H1 onward, inconsistent with a smooth continuation of the pre-ChatGPT trend.

Event study. Figures A10–A15 plot the half-year event study coefficients (Equation 2) for all six age groups, with 2022H1 as the reference period. Table A9 reports a joint test of whether pre-treatment coefficients are collectively zero.¹ The joint test rejects for all age groups (Table A9), reflecting

¹We compute a Wald statistic $W = \hat{\boldsymbol{\delta}}' \hat{V}^{-1} \hat{\boldsymbol{\delta}}$, where $\hat{\boldsymbol{\delta}}$ is the vector of pre-treatment coefficients and \hat{V} is the corresponding block of the clustered variance-covariance matrix. Under the null, $W \sim \chi^2(k)$ where k is the number of pre-treatment periods. Passing or failing a pre-test is neither necessary nor sufficient for parallel trends (Roth, 2022); the employer×month fixed effects provide the primary identification, not the pre-trend test.

Table A8: Employer-level difference-in-differences: employment effect of AI exposure by age group. Dependent variable: $\ln(\text{employment} + 1)$. All specifications include $\text{employer} \times \text{quartile}$ and $\text{employer} \times \text{month}$ fixed effects. Standard errors clustered by $\text{employer} \times \text{quartile}$.

	22–25	26–30	31–34	35–40	41–49	50+
$\hat{\gamma}_1$ (PostRB \times Q4)	−0.012*** (0.002)	−0.006*** (0.002)	−0.001 (0.002)	−0.002 (0.002)	+0.001 (0.001)	+0.005*** (0.001)
$\hat{\gamma}_2$ (PostGPT \times Q4)	−0.010*** (0.002)	+0.000 (0.002)	+0.008*** (0.002)	+0.003** (0.002)	+0.002 (0.001)	+0.010*** (0.001)
N	11,970,426	16,498,560	14,773,590	17,511,078	21,810,984	29,895,762
Emp. \times quartile FE	Yes	Yes	Yes	Yes	Yes	Yes
Emp. \times month FE	Yes	Yes	Yes	Yes	Yes	Yes

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

the large sample size (11–30 million observations per group) rather than economically meaningful pre-trends: pre-period coefficients range from -0.016 to $+0.013$ across all groups, substantially smaller than the post-2024 effects. For ages 22–25 (Figure A10), the post-treatment path shows coefficients near zero through 2023H2 before dropping to -0.029 in 2024H1 and accelerating to -0.055 in 2025H1, consistent with the diffusion timeline of enterprise AI tools. The 26–30 group follows a similar but lagged trajectory, reaching -0.049 by 2025H1. For ages 41–49, coefficients are close to zero throughout, confirming no AI-related employment effect for prime-age workers. A Rambachan and Roth (2023) sensitivity analysis, which constructs honest confidence intervals under plausible violations of parallel trends, is a natural complement to these pre-trend tests; we apply it to the posting event study (Section II.4) and plan to extend it to the employment event study in a revision.

Interpreting pre-trend violations. For ages 41–49 and 50+, the pre-treatment event study coefficients trace an upward path (from negative toward zero), indicating that employment in high-AI occupations was already converging toward low-AI occupations before ChatGPT. If this positive trend continued into the post-treatment period, it biases the DiD coefficient upward. However, the pre-trend slope is modest (approximately 0.002–0.003 per half-year), far too small to mask a negative treatment effect comparable to the 22–25 group.

Table A9: Joint pre-trend test for employer-level event study. H_0 : all pre-treatment coefficients $\hat{\delta}_s = 0$ for $s < 2022H1$.

Age group	χ^2	d.f.	p -value
22–25	383.25	6	0.000
26–30	544.89	6	0.000
31–34	189.85	6	0.000
35–40	169.68	6	0.000
41–49	126.74	6	0.000
50+	179.42	6	0.000

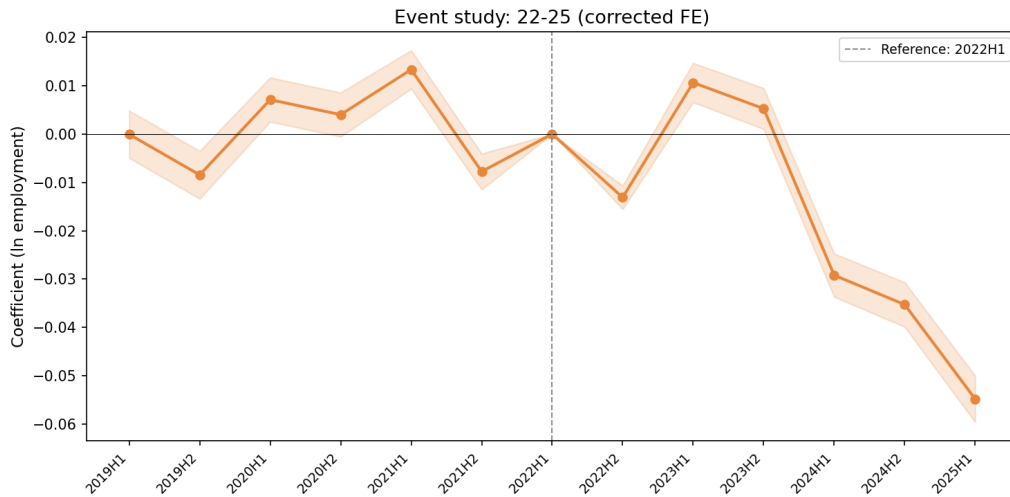


Figure A10: Half-year event study: ages 22–25. Coefficients on $\text{HighQ4} \times \text{half-year period}$ interactions (reference: 2022H1). Employer \times quartile and employer \times month fixed effects. Shaded area: 95% confidence interval. The divergence begins in 2024H1; by 2025H1 the coefficient reaches -0.055 .

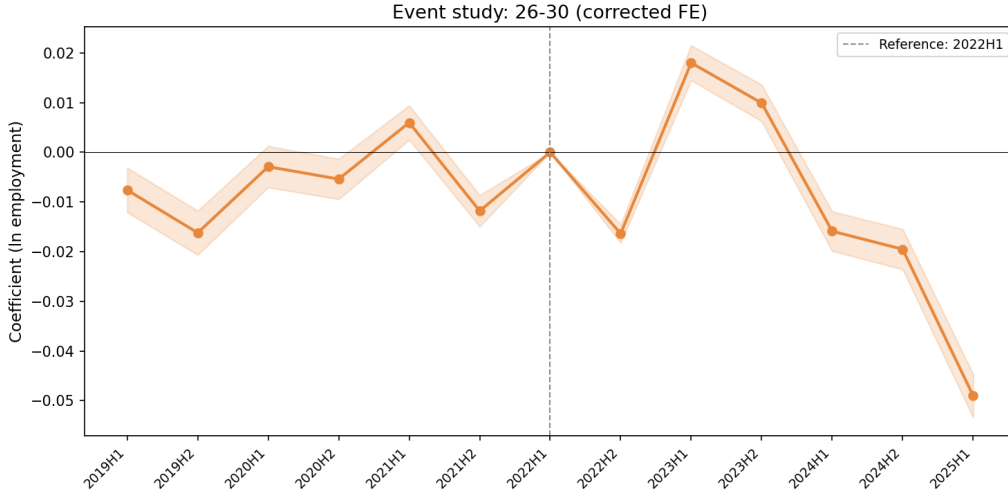


Figure A11: Half-year event study: ages 26–30. Same specification as Figure A10 (employer×quartile and employer×month FE). A similar post-treatment decline emerges from 2024H1, reaching -0.049 by 2025H1.

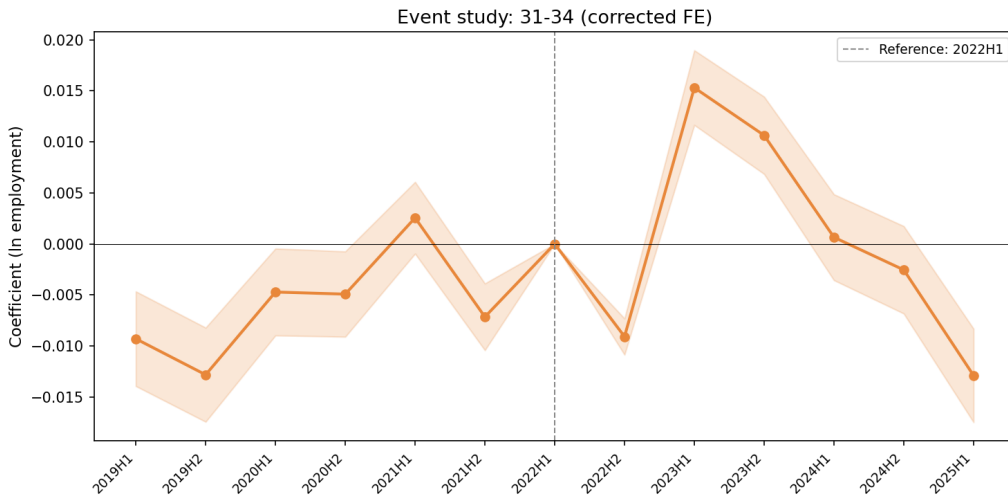


Figure A12: Half-year event study: ages 31–34. Same specification as Figure A10 (employer×quartile and employer×month FE). The average DiD is positive ($+0.008$), but the event study shows a late decline to -0.013 by 2025H1.

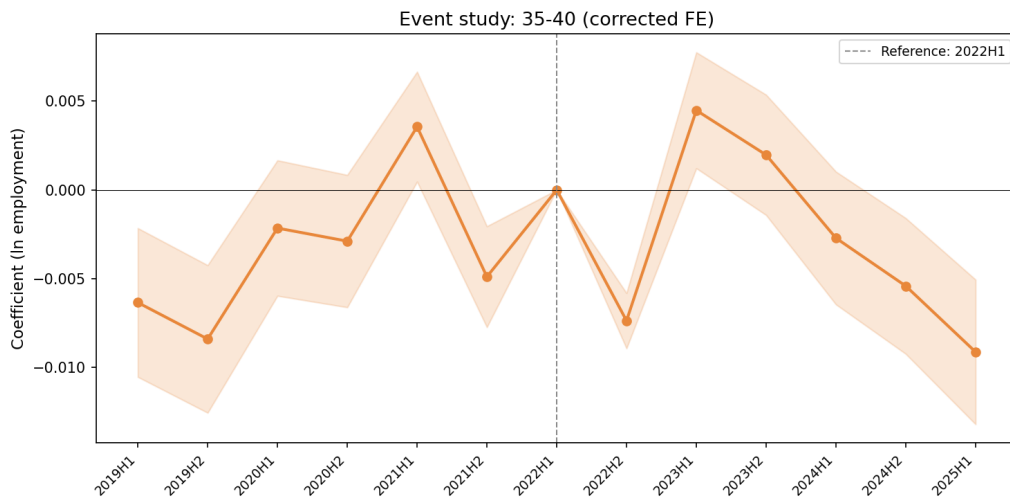


Figure A13: Half-year event study: ages 35–40. Same specification as Figure A10 (employer×quartile and employer×month FE). Small negative effect emerging by 2025H1 (−0.009).

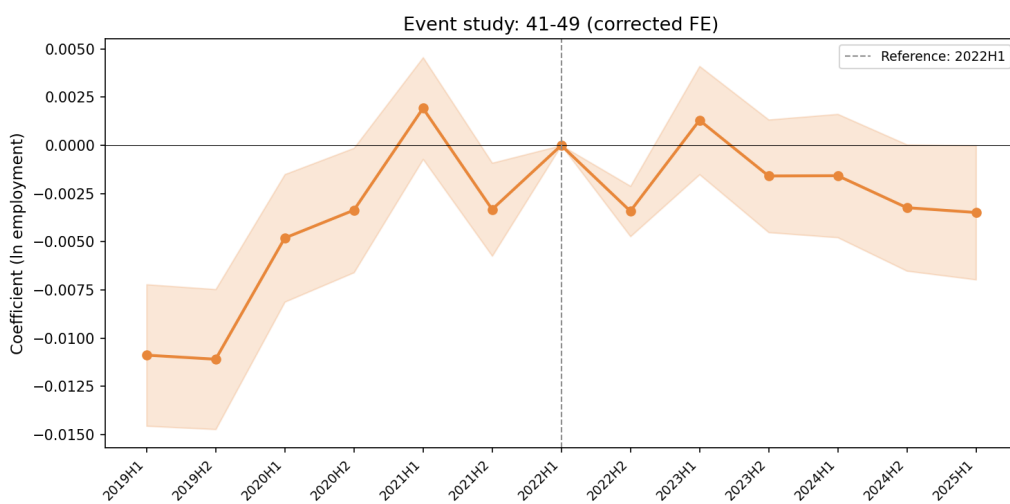


Figure A14: Half-year event study: ages 41–49. Same specification as Figure A10 (employer×quartile and employer×month FE). Coefficients are close to zero throughout, confirming no AI-related employment effect for this age group.

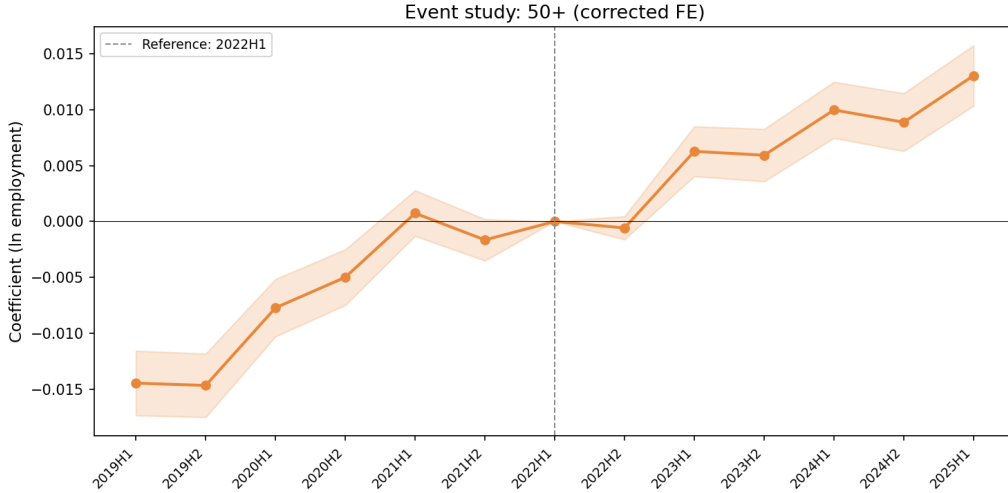


Figure A15: Half-year event study: ages 50+. Same specification as Figure A10 (employer \times quartile and employer \times month FE). The coefficient is positive and rising, reaching +0.013 by 2025H1.

For ages 26–30, the pre-trend is likewise upward, yet the post-treatment path breaks sharply downward: coefficients fall to -0.049 by 2025H1, far exceeding any linear extrapolation of the pre-trend. The average DiD coefficient (+0.000) masks the late-period effect because it pools the mildly positive early post-treatment (2022H2–2023H2) with the sharp late decline (2024H1–2025H1). By the most recent period, the 26–30 effect size approaches that of the 22–25 group, suggesting that AI-related employment displacement is diffusing up the age distribution as adoption deepens. A formal sensitivity analysis along the lines of [Rambachan and Roth \(2023\)](#) could quantify the robustness of these conclusions to alternative trend assumptions.

Cumulative pre-trend bias. Table A10 quantifies the scope for pre-trend contamination of the post-treatment estimates. For each age group, we fit a linear trend through the six pre-treatment event study coefficients (2019H1–2021H2) and extrapolate over the six post-treatment half-years to compute cumulative bias. The pre-trend slope is *positive* for all age groups, including the core 22–25 group (+0.0007/half-year, $R^2 = 0.02$) and the 26–30 group (+0.0012/half-year, $R^2 = 0.09$). Because the post-treatment effects for these groups are negative (-0.055 and -0.049 by 2025H1), linear extrapolation of pre-trends biases *against* the canaries finding: the true effects may be larger

than estimated. The near-zero R^2 values confirm that the pre-period coefficients for 22–25 and 26–30 are essentially noise rather than a systematic trend.

For the 50+ group, the picture is different: the pre-trend slope (+0.003/half-year, $R^2 = 0.88$) extrapolates to a cumulative bias of +0.019, exceeding the observed endpoint (+0.013). The positive 50+ result may therefore partly reflect trend continuation rather than a treatment effect, and should be interpreted with caution.

Table A10: Cumulative pre-trend bias analysis. Pre-trend slope fitted through the six pre-treatment event study coefficients (2019H1–2021H2) via OLS. Cumulative bias = slope \times 6 post-treatment half-years. “Bias/endpoint” gives the ratio of cumulative bias to the 2025H1 coefficient; negative ratios indicate that the pre-trend works *against* the finding.

Age group	Pre-trend slope	R^2	\times periods	Cum. bias	Endpoint	Bias/endpoint
22–25	+0.0007	0.02	$\times 6$	+0.004	−0.055	−7%
26–30	+0.0012	0.09	$\times 6$	+0.007	−0.049	−15%
31–34	+0.0016	0.34	$\times 6$	+0.010	−0.013	−75%
35–40	+0.0012	0.30	$\times 6$	+0.007	−0.009	−80%
41–49	+0.0022	0.70	$\times 6$	+0.013	−0.004	−385%
50+	+0.0032	0.88	$\times 6$	+0.019	+0.013	+149%

III.3. Placebo date test

A direct way to assess whether the post-ChatGPT acceleration is specific to the AI channel is to re-estimate the event study with a placebo treatment date. We set the fake treatment to July 2022, halfway between the Riksbank’s first rate hike (April 2022) and the ChatGPT launch (November 2022), with 2022H1 as the reference period. If the age-gradient effect reflects a pre-existing trend or the Riksbank tightening rather than AI, a comparable divergence should appear at this placebo date.

Figure A16 shows the results. For ages 22–25, the coefficient at the placebo date (2022H2) is −0.013, comparable to the modest average DiD, but the trajectory through 2023 is flat or slightly positive (+0.011 in 2023H1, +0.005 in 2023H2). The sharp acceleration begins only in 2024H1 (−0.029) and deepens to −0.055 by 2025H1. This timing pattern is inconsistent with a smooth continuation of pre-existing trends or a Riksbank-driven effect; instead, it aligns with the enterprise AI adoption timeline. For workers over

50, the upward trajectory is more gradual, consistent with the pre-trend continuation noted in Table A10.



Figure A16: Event study with placebo treatment date: July 2022 (reference period: 2022H1). Employer \times quartile and employer \times month fixed effects. Vertical dashed lines mark the Riksbank’s first rate hike and the ChatGPT launch. The key finding is the flat trajectory through 2023 for ages 22–25, followed by a sharp acceleration in 2024–2025 that is absent at the placebo date.

III.4. Extensive margin: zero-employment cells

Our balanced panel includes zero-employment cells (employer \times quartile \times month observations where no worker of a given age group is employed). These zeros are informative: a rising zero-cell share in Q4 relative to Q1–Q3 indicates that employers are ceasing to employ *any* young workers in high-AI-exposure occupations, the extensive-margin counterpart of the intensive-margin coefficients in Table A8.

Figure A17 plots the zero-cell share separately for Q4 (high AI exposure) and Q1–Q3 (low exposure) by age group. For ages 22–25, the Q4–Q1Q3 gap in zero-cell share widens from approximately 8.9 percentage points in 2022H1 to 13.1 percentage points by 2025H1, indicating that employers are increasingly likely to have no young workers at all in AI-exposed roles. For workers over 50, the gap narrows from 1.2 to 0.7 percentage points, consistent with

Q4 retaining or gaining older workers. The extensive-margin pattern reinforces the event study: the canaries effect operates partly through employers dropping young Q4 employment entirely, not only through reductions in cell size.



Figure A17: Zero-employment cell share in balanced panels: Q4 (high AI exposure) vs Q1-Q3 (low exposure), by age group and half-year period. The widening Q4-Q1Q3 gap for ages 22-25 from 2022H2 onward indicates that employers are increasingly dropping all young workers from AI-exposed roles (extensive margin). For ages 50+, the gap narrows, consistent with retention of experienced workers.

III.5. Gender heterogeneity

Table A11 re-estimates the employer-level DiD (Equation 1) separately for men and women. The negative effect on the youngest workers is approximately twice as large for women as for men: $\hat{\gamma}_2 = -0.016$ ($p < 0.001$) for women aged 22-25, compared with -0.007 ($p < 0.01$) for men. Women aged 26-30 also show a significant negative effect (-0.006 , $p < 0.01$), while men in the same age group show no effect. At the older end, men in high-exposure occupations gain more than women ($+0.016$ vs $+0.001$, the latter not significant). Women aged 41-49 show a significant negative coefficient (-0.009 , $p < 0.001$), unlike men in the same group ($+0.005$, $p < 0.01$). Figure A18 shows the raw employment trajectories by gender.

Table A11: Employer-level DiD by gender and age group. Dependent variable: $\ln(\text{employment} + 1)$. All specifications include $\text{employer} \times \text{quartile}$ and $\text{employer} \times \text{month}$ fixed effects. Standard errors clustered by $\text{employer} \times \text{quartile}$.

	22–25	26–30	31–34	35–40	41–49	50+
<i>Panel A: Men</i>						
$\hat{\gamma}_1$ (PostRB \times Q4)	−0.012*** (0.002)	−0.004 (0.002)	−0.003 (0.002)	+0.000 (0.002)	+0.006*** (0.002)	+0.008*** (0.001)
$\hat{\gamma}_2$ (PostGPT \times Q4)	−0.007*** (0.003)	−0.002 (0.002)	+0.007*** (0.002)	+0.003 (0.002)	+0.005*** (0.002)	+0.016*** (0.001)
<i>N</i>	6,401,850	9,373,104	8,448,726	10,064,574	12,338,664	16,984,266
<i>Panel B: Women</i>						
$\hat{\gamma}_1$ (PostRB \times Q4)	−0.006*** (0.002)	−0.003* (0.002)	+0.006*** (0.002)	+0.001 (0.002)	−0.001 (0.002)	+0.004** (0.001)
$\hat{\gamma}_2$ (PostGPT \times Q4)	−0.016*** (0.002)	−0.006*** (0.002)	+0.004* (0.002)	+0.002 (0.002)	−0.009*** (0.002)	+0.001 (0.002)
<i>N</i>	7,412,574	9,506,640	7,746,258	8,826,948	10,863,138	14,592,396
Emp. \times quartile FE	Yes	Yes	Yes	Yes	Yes	Yes
Emp. \times month FE	Yes	Yes	Yes	Yes	Yes	Yes

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

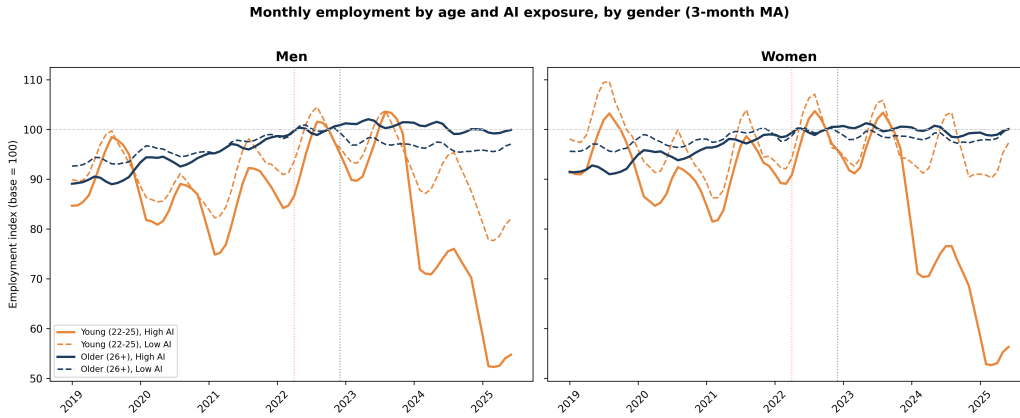


Figure A18: Monthly employment by age group, AI exposure, and gender (3-month moving average). AGI register data indexed to October 2022 = 100. Young = 22–25; Older = 26+. Solid lines = high AI exposure (Q4); dashed = low (Q1–Q3). The decline among young women in high-exposure occupations is steeper than for young men.

Decomposing the gender gap: occupational composition vs within-occupation effects. The gender differences in Table A11 could reflect either (i) occupational composition – men and women sort into different occupations within the top AI exposure quartile – or (ii) within-occupation differences in how AI exposure affects employment by gender. To separate these channels, we apply an Oaxaca-Blinder-style reweighting decomposition: we re-estimate the women’s regressions using inverse probability weights (IPW) that adjust women’s occupational distribution within Q4 to match men’s. If the gender gap narrows after reweighting, the between-occupation (composition) component is a contributing factor; if the gap persists, the within-occupation (coefficient) component dominates.

Table A12 reports the results. For ages 22–25, reweighting reduces the women’s coefficient from -0.016 to -0.013 ($p < 0.001$), narrowing the gender gap from 0.009 to 0.006. Occupational composition accounts for approximately one-third of the gender gap; the remaining two-thirds reflects within-occupation differences, indicating that young women in the *same* AI-exposed occupations as young men experience larger employment declines.

The composition channel is more consequential at older ages. For women aged 41–49, the significant negative coefficient (-0.009 , $p < 0.001$) disappears entirely after reweighting ($+0.002$, $p = 0.42$). The effect for middle-aged women is thus fully attributable to their concentration in specific Q4 occupations – likely administrative and customer service roles – that are disproportionately affected, not to differential treatment within occupations. Similarly, for women over 50, the insignificant unweighted coefficient ($+0.001$, $p = 0.52$) becomes a significant gain after reweighting ($+0.011$, $p < 0.001$), comparable in magnitude to men’s gain ($+0.016$). Middle-aged and older women’s occupational sorting into the hardest-hit Q4 occupations masks an underlying pattern that, conditional on occupation, resembles men’s.

III.6. Spotlight occupations: monthly employment by age (MONA)

Brynjolfsson et al. (2025) highlight specific “spotlight” occupations (software developers and customer service agents) where AI adoption is particularly visible. A forecast report independently documents declining employment in administration and customer service in Sweden, noting that “unemployment among white-collar workers has risen faster than among blue-collar workers” (Brännemo, 2025). Using monthly AGI (employer declaration) register data from SCB’s secure environment (MONA), we plot employment

Table A12: Decomposing the gender gap: women’s DiD coefficients, unweighted vs IPW-reweighted to men’s occupational distribution within Q4. All specifications include employer×quartile and employer×month fixed effects.

	Men (unweighted)	Women (unweighted)	Women (reweighted)
22–25	−0.007*** (0.003)	−0.016*** (0.002)	−0.013*** (0.002)
26–30	−0.002 (0.002)	−0.006*** (0.002)	−0.005** (0.002)
31–34	+0.007*** (0.002)	+0.004* (0.002)	+0.006* (0.003)
35–40	+0.003 (0.002)	+0.002 (0.002)	+0.001 (0.003)
41–49	+0.005*** (0.002)	−0.009*** (0.002)	+0.002 (0.002)
50+	+0.016*** (0.001)	+0.001 (0.002)	+0.011*** (0.002)

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. $\hat{\gamma}_2$ (PostGPT × Q4) shown. Standard errors in parentheses.

trajectories by fine age band (22–25, 26–30, 31–34, 35–40, 41–49, 50+) for these occupations, indexed to October 2022 (just before ChatGPT).

The economy-wide employment figure by age group and AI exposure appears as Figure 2 in the main paper. Below we show occupation-specific spotlight analyses for four high-exposure occupations that span the augmentation–automation spectrum.

Software developers (SSYK 2512) represent a setting where AI primarily augments experienced workers: senior developers leverage code-generation tools to become more productive, potentially reducing demand for junior coding tasks. Payroll administrators (SSYK 4112) and receptionists (SSYK 4225) represent occupations where routine entry-level tasks – data entry, scheduling, basic enquiries – are more directly automatable. Customer service agents (SSYK 4222) lie between these poles. That all four occupations display the same age gradient – young workers declining, older workers stable or rising – suggests that both augmentation and automation channels reduce demand for entry-level workers, albeit through different mechanisms: augmentation by raising the productivity of experienced workers, automation by eliminating codified tasks disproportionately performed by juniors.

The augmentation–automation spectrum in these spotlight occupations

connects to the gender gap documented in Section III.5. Three of the four occupations are female-dominated: payroll administrators (82% female), customer service agents and receptionists (66% female), while software developers are 76% male (SCB Yrkesregistret 2024, 3-digit SSYK). The Oaxaca-Blinder decomposition shows that reweighting women to men’s occupational distribution within Q4 narrows the gender gap by one-third, consistent with women’s concentration in the administrative and service occupations where age-gradient effects are largest.

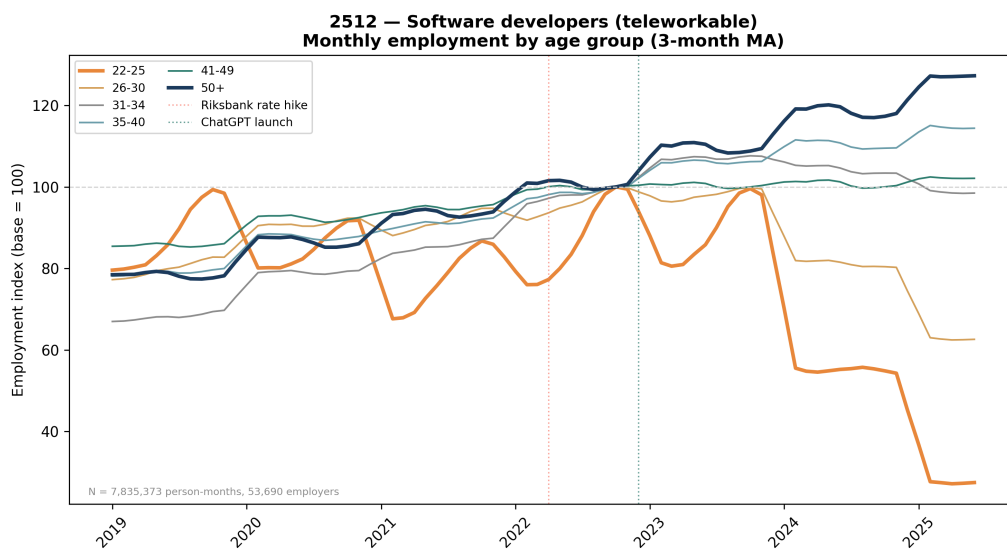


Figure A19: Monthly employment by age group: Software Developers (SSYK 2512). AGI register data, indexed to base month = 100. Vertical lines mark the Riksbank’s first rate hike (April 2022) and ChatGPT launch (November 2022). The youngest cohort (22–25) shows the steepest decline, consistent with the economy-wide canaries effect in Figure 2 of the main paper.

III.7. Percentage change in employment

The indexed employment series in Figure 2 of the main paper shows levels relative to October 2022 (= 100). To make the magnitude of the divergence immediately transparent, Figures A23–A24 recast the same data as percentage changes from October 2022, following the presentation in Brynjolfsson et al. (2025).

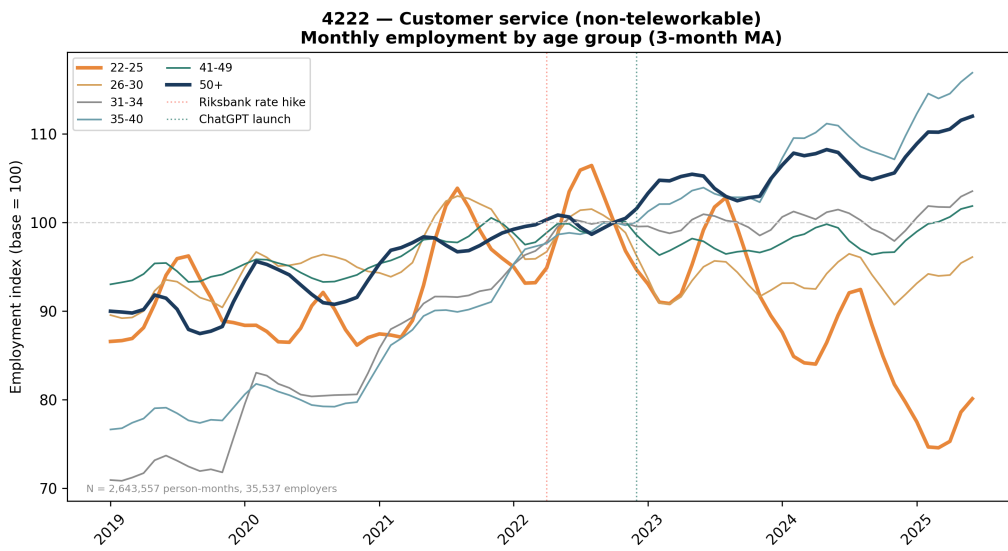


Figure A20: Monthly employment by age group: Customer Service Agents (SSYK 4221, 4222, 5230). AGI register data, indexed to base month = 100. Vertical lines mark the Riksbank’s first rate hike (April 2022) and ChatGPT launch (November 2022). The pattern is noisier than for software developers, potentially reflecting smaller cell sizes and more heterogeneous task content.

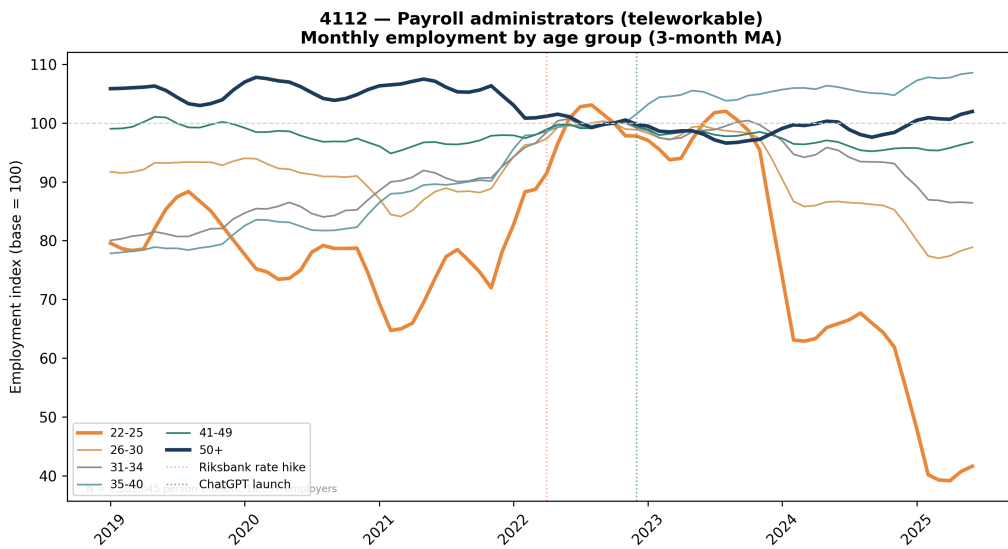


Figure A21: Monthly employment by age group: Payroll Administrators (SSYK 4112). AGI register data, indexed to base month = 100. Vertical lines mark the Riksbank's first rate hike (April 2022) and ChatGPT launch (November 2022). Payroll administrators rank at the 96th percentile of DAIOE genAI exposure. The 22–25 group drops to approximately 60% of the base by early 2025.

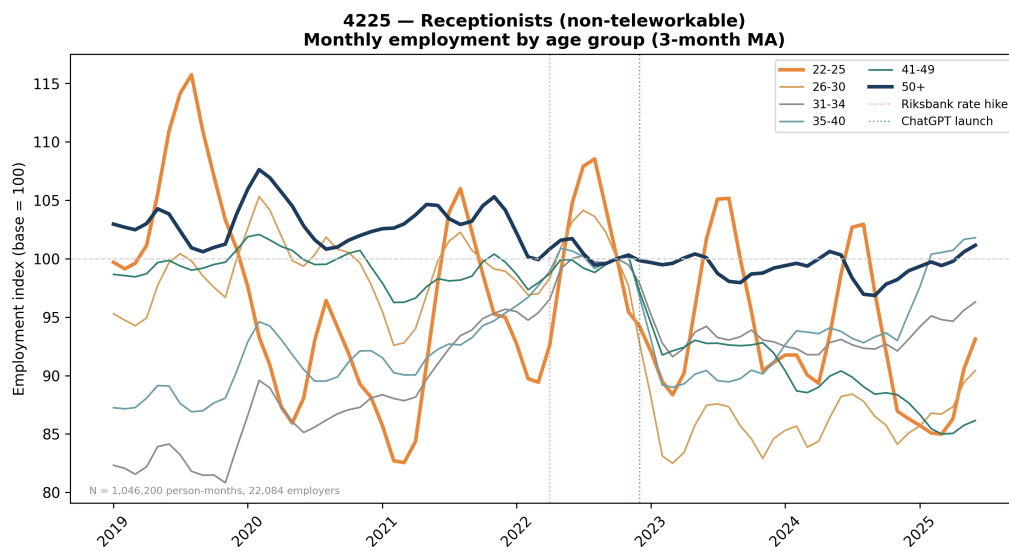


Figure A22: Monthly employment by age group: Receptionists (SSYK 4225). AGI register data, indexed to base month = 100. Vertical lines mark the Riksbank’s first rate hike (April 2022) and ChatGPT launch (November 2022). Receptionists rank at the 76th percentile of DAIOE genAI exposure. The pattern is noisier than for software developers or payroll administrators, likely reflecting smaller cell sizes.

Figure A23 isolates the trajectory of 22–25 year olds in top-quartile AI-exposed occupations: employment fell 44 per cent from October 2022 to early 2025. The strong seasonal pattern visible in the pre-period (summer employment dips) is dwarfed by the post-ChatGPT trend. Figure A24 adds the three comparison groups. Older workers (26+) in both high and low AI exposure remain near zero throughout; young workers (22–25) in low-exposure occupations decline modestly (−11%). The four-way comparison makes both the age gradient and the exposure gradient visible in a single figure.

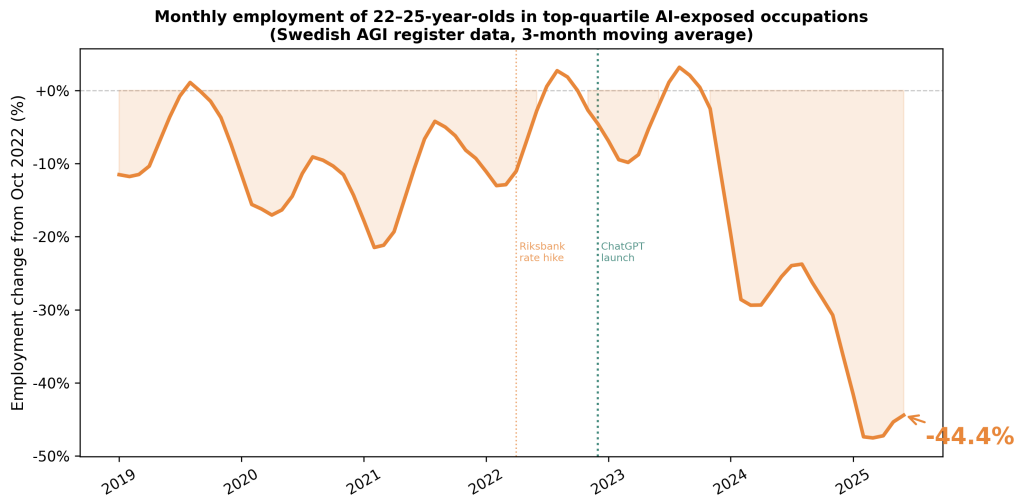


Figure A23: Employment change for 22–25 year olds in top-quartile AI-exposed occupations. Percentage change from October 2022, 3-month centred moving average. AGI register data.

III.8. Comparison with related studies

Table A13 compares our study with the emerging literature testing whether generative AI exposure affects employment or job postings differentially by worker age or seniority. Studies vary along several dimensions: data source (payroll, register, job postings), geographic coverage, AI exposure measure, identification strategy, and whether age or seniority heterogeneity is tested directly.

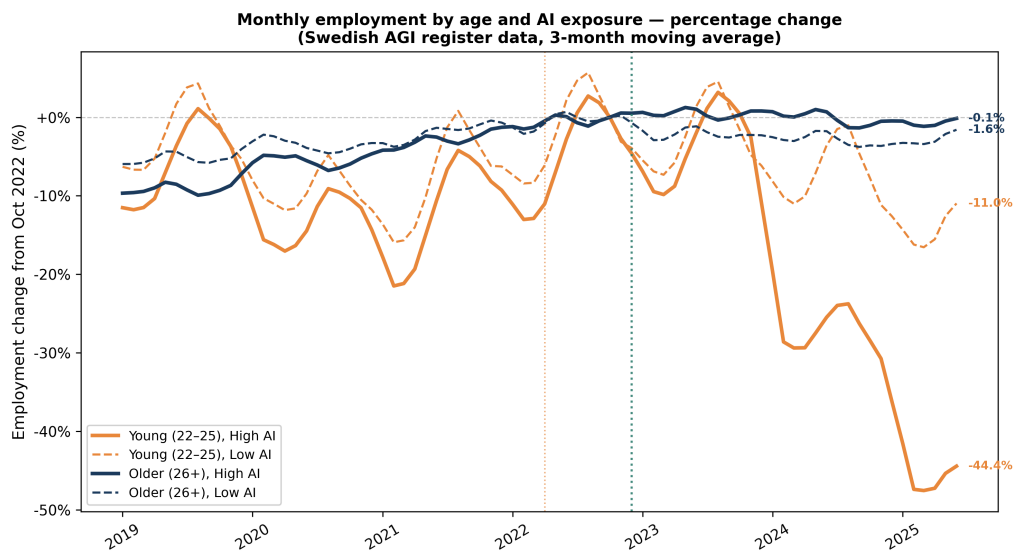


Figure A24: Employment change by age group and AI exposure. Percentage change from October 2022, 3-month centred moving average. AGI register data. Young = 22–25; Older = 26+. High AI = top quartile of DAIOE genAI exposure; Low AI = bottom three quartiles.

Three patterns emerge. First, Anglo-Saxon studies using job posting or payroll data consistently find negative effects on entry-level or junior workers (Brynjolfsson et al., 2025; Hosseini Maasoum and Lichtinger, 2025; Liu et al., 2025; Klein Teeselink, 2025). Second, Nordic register-based studies show mixed results: our Swedish evidence finds a significant age gradient at the employer level, while Finnish analyses report null effects both at the occupation level (Kauhanen and Rouvinen, 2025) and – notably – under the same firm-level Poisson design used by Brynjolfsson et al. (2025) (Kauhanen and Rouvinen, 2026), and the Danish intensive-margin analysis finds precise nulls on earnings and hours (Humlum and Vestergaard, 2025). Third, sceptical contributions challenge the emerging consensus: Eckhardt and Goldschlag (2025) find that unemployment rose *less* in AI-exposed occupations using CPS data, and Iscenko and Curto Millet (2026) argue the US pattern reflects monetary policy rather than AI. Our Riksbank timing test directly addresses the latter hypothesis: the posting decline aligns with the rate hike, but the within-employer age gradient in employment emerges only after ChatGPT.

Our study is the first to combine both margins (postings and employment) with a natural timing test separating macroeconomic and AI channels, using population-register data with employer-level identification. For broader surveys of this emerging literature, see Chandar (2025) and Kolko (2026).

Table A13: Comparison of studies testing AI exposure effects on employment or job postings. FE = fixed effects; ES = event study; DiD = difference-in-differences; SDiD = synthetic DiD.

Study	Country	Data	Method	FE structure	Age tested?	Key result
Brynjolfsson et al. (2025)	US	ADP payroll, 3.5–5M workers/month	Poisson ES	Firm×quintile, firm×month	Yes	−16%, ages 22–25
Hosseini Maasoum and Lichtinger (2025)	US	Resumes + postings, 62M workers	Firm DiD	Firm FE	Yes	Junior decline, senior stable
Klein Teeseling (2025)	UK	Lightcast + LinkedIn	DiD	Firm/occupation	Yes	Junior roles −5.8%
Liu et al. (2025)	US	Lightcast, 285M postings	DiD + ES	Occupation + time	Yes	Entry-level postings −18% by year 3
Storm et al. (2025)	Germany	IAB admin + vacancies	IV	Individual + occ×year	Indirect (skill)	Experts gain, non-experts lose
Eckhardt and Goldschlag (2025)	US	CPS	Descriptive	CPS controls	No	Unemployment rose <i>less</i> in high-exposure occ.
Ischenko and Curto Millet (2026)	US	Lightcast + CPS	Descriptive	—	Yes	Argues null; attributes pattern to monetary policy
Massenkoff and McCrory (2026)	US	CPS + Claude usage	Descriptive	CPS controls	Suggestive	Job finding rate in exposed occ. −14%, ages 22–25
Kauhanen and Rouvinen (2025)	Finland	Income register, population	SDiD	Occupation-level	No	Null on employment and wages
Kauhanen and Rouvinen (2026)	Finland	Income register, population	Poisson ES	Firm×quintile, firm×month	Yes	Null; if anything positive for young
Humlum and Vestergaard (2025)	Denmark	Admin register + survey	Pre/post	Occupation-level	No	Null on earnings and hours
This paper	Sweden	AGI register + Platsbanken, population	DiD + ES	Emp×quartile, emp×month	Yes	−5.5% by 2025H1 (ES), ages 22–25; postings null

Part IV: Data documentation

IV.1. Data sources and access

Table A14 summarises the data sources used in this paper. Detailed documentation for each source follows.

Table A14: Data sources.

Source	Period	Frequency	Unit	Key variables
Platsbanken (AF)	2020–2026	Monthly	Job ad	SSYK 4-digit, vacancies, municipality, employer
AGI register (SCB/MONA)	2019–2025	Monthly	Worker–employer	Employment, SSYK 4-digit, age, gender, employer ID
DAIOE	2023	Cross-section	SSYK 4-digit	GenAI exposure percentile rank
OMXS30	2020–2026	Daily → monthly	Index	Closing price
Riksbank policy rate	2020–2026	Event dates	—	Policy rate level

IV.1.1. *Platsbanken*

Job advertisement data from the Swedish Public Employment Service (Arbetsförmedlingen, AF), published under CC0 licence. Each record contains: ad identifier, publication date, SSYK 2012 four-digit occupation code, number of vacancies, municipality code, employer name, and source type. For 2020–2025 we use the historical bulk archive (<https://data.jobtechdev.se/annonser/historiska/>); for January–February 2026 we supplement with the real-time JobStream API (<https://jobstream.api.jobtechdev.se>), which provides the same fields for currently published ads. The two sources are deduplicated on ad identifier before aggregation. In recent years, approximately 700,000 ads have been published annually on Platsbanken.

Coverage and representativeness. Job postings data are increasingly used in labour market research (e.g. Engberg et al., 2024). Platsbanken is the largest recruitment site in Sweden, covering the near-universe of advertised job vacancies. Several considerations bear on the representativeness of these data.

First, most vacancies in Sweden are filled through channels other than formal advertising, such as personal contacts and direct approaches. However,

the share of vacancies filled via AF has been stable or slightly increasing over time: in 2006, 9% of already-employed job changers and 13% of unemployed individuals who found work did so via AF; by 2023 these shares had risen by 1 and 2 percentage points, respectively.

Second, private and local-government employers were previously required by law (SFS 1976:157) to post their vacancies at AF, but the obligation was repealed in 2007 (SFS 2007:396); only central government retains its own posting requirement (SFS 1984:819). [Cronert \(2022\)](#) finds that the regulatory change reduced posting propensity by up to about 18 per cent – a relatively modest effect, likely because the earlier mandate was in practice not enforced.

Third, recruitment advertising has shifted toward digital platforms. Besides Platsbanken, vacancies are posted on LinkedIn, Monster, and various sector-specific sites. That Platsbanken retains its position likely reflects established practice and the fact that posting is free; AF has also facilitated re-distribution of postings to other platforms. Nevertheless, for certain industries and occupations, Platsbanken may undercount the number of vacancies.

Fourth, existing employees may be redeployed internally to AI-related tasks. Since AI competence is sought-after, organisations may rely relatively more on internal rather than external recruitment, meaning that job postings potentially understate AI adoption.

For our analysis, the key question is whether any measurement gap varies systematically with AI exposure. To the extent that high-AI-exposure occupations (e.g., white-collar professional roles) are more likely to also be advertised on competing platforms, our analysis would undercount postings in Q4 relative to Q1, biasing us *toward* finding a negative AI effect. The null result for $\hat{\beta}_2$ in the main paper is therefore conservative with respect to platform migration.

IV.1.2. DAIOE

The Dynamic AI Occupational Exposure index maps AI benchmark performance to occupational task content. The genAI variant focuses on capabilities relevant to large language models and image generation. Publicly available; see [Engberg et al. \(2024\)](#).

IV.1.3. OMXS30

Stockholm OMX 30 daily closing prices from Yahoo Finance (ticker: `^OMX`).

IV.1.4. Riksbank policy rate

Manually compiled from Riksbank press releases. Key dates verified against <https://riksbank.se>. Figure A8 plots the full tightening cycle.

IV.1.5. Administrative employment data (SCB/MONA)

The employment analysis uses monthly employer declaration (AGI, *arbetsgivardeklaration*) register data from Statistics Sweden (SCB). These data are drawn from the ORU-MICRO-AI employer–employee database of Örebro University, which links several administrative registers. The variables from the database that are used in this paper are: encrypted person identifier, SSYK 2012 four-digit occupation code (from the 2023 register year), age, gender, employer identifier, and monthly employment status.

Source and coverage. The AGI register originates from mandatory monthly employer declarations to the Swedish Tax Agency (*Skatteverket*), introduced in January 2019. Every employer making PAYE (*A-skatt*) salary payments must report each individual payment, making the AGI a population-level register of formal employment. Unlike survey-based sources, the AGI has no sampling error and covers the full universe of employer–employee matches, including part-time and temporary workers. Since the unit of observation is the employer–employee relationship, an individual holding positions at multiple employers contributes one record per employer per month; duplicates within the same employer–occupation–month cell are removed before aggregation. Self-employed individuals operating under F-skatt (sole traders, platform workers) do not appear in the AGI; a shift toward such arrangements among young workers would reduce measured employment but would not differentially affect AI-exposed occupations. We restrict the sample to employers with at least five workers, because very small employers generate noisy occupation×age cells and their hiring patterns are too idiosyncratic to identify within-employer composition shifts across AI exposure quartiles.

Occupation coding. SSYK 2012 four-digit occupation codes are employer-reported and recorded in a separate annual occupation register (*Yrkesregistret*). We assign each worker the occupation from the most recent available register year using a cascading lookup: the 2023 register is consulted first; if no code is found, the 2022 and then the 2021 register are used. For data years prior to 2023, workers are matched to the corresponding year’s own register. Since occupational transitions at the 4-digit level are slow (annual

transition rates are approximately 5–8% for most age groups), this assignment introduces minimal measurement error. A graduate entering a high-AI occupation in 2024 may be coded to their 2023 student job, which would attenuate rather than inflate the measured age gradient. Conversely, if workers increasingly move into high-exposure occupations over time, the Q4 employment count would be mechanically inflated, again working against our finding of a relative decline in Q4 employment among young workers.

Stability and comparability. The AGI data source has been stable since its introduction in January 2019, with no changes to the reporting requirements or the SSYK 2012 classification during the study period (2019–2025). This contrasts with the annual employment registers, where SCB replaced the RAMS methodology with BAS from reference year 2022, introducing a methodological break at our treatment timing. Our use of monthly AGI data avoids this break entirely.

Non-match rates and sample attrition. Matching individuals to SSYK 4-digit occupations yields non-match rates of 11 per cent for ages 22–25 and 7 per cent for ages 26 and above. The higher non-match rate among younger workers likely reflects short-term contracts and recent hires for which employer declarations have not yet been updated with occupation codes. Unmatched individuals are excluded from the analysis. We find no evidence that the non-match rate changes differentially across high- and low-AI-exposure occupations after the treatment date (see Section III.2).

The overall SSYK non-match rate is approximately 10 per cent through 2023, rising to 15 per cent in 2024 and 20 per cent in 2025. The increase reflects the end of LISA coverage at 2023: new labour market entrants from 2024 onward lack a SSYK code unless their employer has already reported one. Importantly, our core results are strongest in 2023, when occupation coverage is most complete. Among young entrants first observed in 2023, only 13 per cent appear in top-quartile (Q4) occupations, compared with 28 per cent in the working population, indicating that attrition is under-represented in high-exposure occupations. Missing individuals have slightly better average labour market outcomes than matched individuals, but this pattern is stable over time, suggesting no differential selection correlated with the treatment.

Access. The administrative micro-level data in the ORU-MICRO-AI database are from several registers of Statistics Sweden, and were obtained after ap-

proval from the Swedish Ethical Review Authority (approvals 2021-05040, 2022-03330-02, 2024-01714-0, 2025-04205-02). Researchers and their assistants may access these and similar Swedish micro-level data, subject to approval by SCB. Access is provided through the secure internet-based Microdata Online Access (MONA) system; microdata never leave SCB's servers. Foreign-based researchers may visit a Swedish institution with access or cooperate with researchers in Sweden. Information on the application process is available at <https://www.scb.se/en/services/ordering-data-and-statistics/ordering-microdata/> or by contacting SCB at mona@scb.se. Researchers wishing to access our specific data for replication purposes should contact the corresponding author for guidance regarding project approval from the Swedish Ethical Review Authority and SCB.

All replication code for both public and restricted data is available at <https://github.com/Magnus-L/canaries-sweden>. The authors commit to preserving data access and providing reasonable assistance to replication requests for at least five years following publication.

References

- Brännemo, T., 2025. Unionen om konjunkturen – hösten 2025. Konjunkturprognos. Unionen. Stockholm. URL: https://www.unionen.se/sites/default/files/files/Unionen-om-konjunkturen_host-2025.pdf.
- Brynjolfsson, E., Chandar, B., Chen, R., 2025. Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence. Working Paper. Stanford Digital Economy Lab.
- Callaway, B., Goodman-Bacon, A., Sant'Anna, P.H.C., 2024. Difference-in-Differences with a Continuous Treatment. Working Paper 32117. National Bureau of Economic Research.
- Chandar, B., 2025. AI and labor markets: What we know and don't know. Stanford Digital Economy Lab. URL: <https://digitaleconomy.stanford.edu/news/ai-and-labor-markets-what-we-know-and-dont-know/>. 10 October 2025, updated 20 January 2026.
- Chen, J., Roth, J., 2024. Logs with zeros? Some problems and solutions. The Quarterly Journal of Economics 139, 891–936. doi:[10.1093/qje/qjad054](https://doi.org/10.1093/qje/qjad054).

- Cronert, A., 2022. When the paper tiger bites: Evidence of compliance with unenforced regulation among employers in Sweden. *Regulation & Governance* 16, 1141–1159.
- Dingel, J.I., Neiman, B., 2020. How many jobs can be done at home? *Journal of Public Economics* 189, 104235.
- Eckhardt, S., Goldschlag, N., 2025. AI and Jobs: The Final Word (Until the Next One). The American Worker Project Working Paper. Economic Innovation Group.
- Eloundou, T., Manning, S., Mishkin, P., Rock, D., 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science* 384, 1306–1308. doi:[10.1126/science.adj0998](https://doi.org/10.1126/science.adj0998).
- Engberg, E., Görg, H., Lodefalk, M., Javed, F., Längkvist, M., Monteiro, N.P., Kyvik Nordås, H., Pulito, G., Schroeder, S., Tang, A., 2024. AI Unboxed and Jobs: A Novel Measure and Firm-Level Evidence from Three Countries. Discussion Paper 16717. IZA.
- Hosseini Maasoum, S.M., Lichtinger, G., 2025. Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Resume and Job Posting Data. Working Paper. Harvard University. SSRN 5425555.
- Humlum, A., Vestergaard, E., 2025. Large Language Models, Small Labor Market Effects. Working Paper 33777. National Bureau of Economic Research.
- Ischenko, Z., Curto Millet, F., 2026. Looking for the Ladder: Is AI Impacting Entry-Level Jobs? The American Worker Project Working Paper. Economic Innovation Group.
- Kauhanen, A., Rouvinen, P., 2025. Assessing early labour market effects of generative AI: Evidence from population data. *Applied Economics Letters* doi:[10.1080/13504851.2025.2513973](https://doi.org/10.1080/13504851.2025.2513973).
- Kauhanen, A., Rouvinen, P., 2026. AI Has Not Impacted the Youth Labor Market in Finland. Working Paper 135. ETLA Economic Research.
- Klein Teeselink, B., 2025. Generative AI and Labor Market Outcomes: Evidence from the United Kingdom. Working Paper. SSRN 5516798.

- Kolko, J., 2026. Research on AI and the labor market is still in the first inning. The Hamilton Project, Brookings Institution. URL: <https://www.hamiltonproject.org/publication/post/research-on-ai-and-the-labor-market-is-still-in-the-first-inning/>. 10 March 2026.
- Liu, Y., Wang, H., Yu, S., 2025. Labor Demand in the Age of Generative AI: Early Evidence from the U.S. Job Posting Data. Policy Research Working Paper 11263. World Bank.
- Massenkoff, M., McCrory, P., 2026. Labor market impacts of AI: A new measure and early evidence. Anthropic Research. URL: <https://www.anthropic.com/research/labor-market-impacts>.
- Plikt- och prövningsverket, 2025. In English. URL: <https://www.pliktverket.se/om-myndigheten/in-english>. accessed 2 March 2026.
- Rambachan, A., Roth, J., 2023. A more credible approach to parallel trends. *Review of Economic Studies* 90, 2555–2591.
- Roth, J., 2022. Pretest with caution: Event-study estimates after testing for parallel trends. *American Economic Review: Insights* 4, 305–322. doi:10.1257/aeri.20210236.
- Roth, J., Sant’Anna, P.H.C., Bilinski, A., Poe, J., 2023. What’s trending in difference-in-differences? A synthesis of the recent econometrics literature. *Journal of Econometrics* 235, 2218–2244. doi:10.1016/j.jeconom.2023.03.008.
- Storm, E., Gonschor, M., Schmidt, M.J., 2025. AI in Demand: How Expertise Shapes its (Early) Impact on Workers. Working Paper 61. Institute for Advanced Studies (IHS).
- Zens, G., Böck, M., Zörner, T.O., 2020. The heterogeneous impact of monetary policy on the US labor market. *Journal of Economic Dynamics and Control* 119, 103989.