



AI Exposure in the Swedish context

A comparison of AI-Induced labor market developments in Swedish economic sectors

Jesper Lindbom

Degree project/Independent project • 15 credits

Swedish University of Agricultural Sciences, SLU

Faculty of Natural Resources and Agricultural Sciences/Department of Economics

Political Science – Sustainable Development

Degree project/SLU, Department of Economics 1750 • ISSN 1401-4084

Uppsala 2026



AI Exposure in the Swedish context: A comparison of AI-Induced labor market developments in Swedish economic sectors

Jesper Lindbom

Supervisor: Jonathan Stråle, Swedish University of Agricultural Sciences, Department of Economics
Examiner: Shon Ferguson, Swedish University of Agricultural Sciences, Department of Economics

Credits: 15 credits
Level: First cycle, G2E
Course title: Independent project in Economics, G2E
Course code: EX0903
Programme/education: Political Science – Sustainable Development (BSc)
Course coordinating dept: Department of Economics
Place of publication: Uppsala
Title of series: Degree project/SLU, Department of Economics
Part number: 1750
ISSN: 1401-4084

Swedish University of Agricultural Sciences
Faculty of Natural Resources and Agricultural Sciences (NJ)
Department of Economics

Abstract

This study examines the relationship between generative artificial intelligence exposure and labor market outcomes across Swedish economic sectors. Using AI exposure measures developed by Felten et al. (2021) and converted from the North American Industry Classification System (NAICS) to the Swedish Standard Industrial Classification (SNI), the study identifies sectors with relatively high and low exposure to generative AI technologies. A difference-in-differences framework combined with an event study specification is employed to analyse the effects of AI exposure on hourly wages and employment levels using quarterly data on the sector level.

The results indicate that highly AI-exposed sectors experienced relatively weaker wage developments during the post-treatment period compared to less exposed sectors. The estimated wage effects remain statistically significant across multiple model specifications, although the magnitude decreases following the inclusion of additional control variables. The analysis identifies more limited and less consistent effects on employment levels. The event-study estimates provide partial support for the parallel trends assumption while suggesting that treatment effects evolve gradually over time.

Keywords: Generative artificial intelligence, Artificial intelligence exposure, Labor market, Wages, Employment, Event study, Sweden

Table of contents

Abstract.....	4
List of tables	6
List of figures.....	7
Abbreviations	8
1. Introduction	9
2. Background	12
2.1 Generative AI	12
2.2 AI exposure scores.....	12
2.3 Industry classification systems.....	13
3. Methodology	15
3.1 Data	15
3.2 Industry classification conversion.....	16
3.3 Empirical strategy	17
3.4 Regression equations	18
3.5 Control variables	19
3.6 Graphical structure	20
4. Results	21
5. Discussion	25
5.1 Interpretations	25
5.2 Results in relation to existing literature	26
5.3 Limitations	27
5.4 Methodological improvements	28
6. Conclusion	30
References	31

List of tables

Table 1. Descriptive statistics of dependent and control variables. **Fel! Bokmärket är inte definierat.**

Table 2. Regression table, hourly wage.....23

Table 3. Regression table, employment level.....24

List of figures

Figure 1. DiD event study graph, hourly wage.....	21
Figure 2. DiD event study graph, employment level.	22

Abbreviations

AI	Artificial Intelligence
AIIE	Artificial Intelligence Industry Exposure
AIOE	Artificial Intelligence Occupational Exposure
BLS	Bureau of Labor Statistics
DiD	Difference-in-differences
EFF	Electronic Frontier Foundation
LLM	Large Language Model
NACE	Statistical Classification of Economic Activities
NAICS	North American Industry Classification System
SCB	Statistics Sweden
SNI	Swedish Standard Industrial Classification
O*NET	Occupational Information Network
SUTVA	Stable Unit Treatment Value Assumption

1. Introduction

Since its emergence, generative artificial intelligence (AI) has garnered significant attention concerning its potential for economic impact and to effectivize the modern workplace. This interest in generative AI has contributed to very rapid adoption in certain industries, as workers and employers increasingly incorporate such tools into the workday ([Bick et al., 2024](#)).

Historically, new technologies automate specific tasks within occupations, making workers shift toward complementary tasks. This supports that technology typically restructures tasks rather than eliminating entire occupations ([Autor, 2015](#)). Implementation of generative AI in the early stages shows potential for increasing worker productivity and can allow employees to make up for eventual work inexperience, currently augmenting workflows rather than automating them ([Brynjolfsson et al., 2025](#)). Research consistently highlights that the workplace performance of generative AI programs is dependent on the nature of tasks to which they are applied. Given the proficiencies of generative AI, these have shown to mainly be tasks that are repetitive, language-intensive, and partly structured that can be reviewed and adjusted by humans ([Eloundou et al., 2023](#)).

As the task composition of jobs is a main determinant of AI performance, it becomes intuitive that specific jobs and industries are more likely to experience implementations of artificial intelligence. In order to quantify this breadth of exposure, the scholars Edward Felten, Manav Raj, and Robert Seamans calculate artificial intelligence exposure scores for all United States industries (AIIE) listed by the U.S. Bureau of Labor Statistics. The scores describe the degree to which tasks and abilities prevalent in each respective industry overlap with capabilities where AI is growing more proficient ([Felten et al., 2021](#)).

Previous research has produced mixed evidence regarding the labor market effects of artificial intelligence and automation. Studies of industrial robots generally find negative effects on wages and employment in exposed labor markets ([Acemoglu & Restrepo, 2020](#)). [Acemoglu and Restrepo \(2019\)](#) argue that automation can both replace and create tasks, making its overall labor market effects uncertain. Experimental studies suggest that generative AI can substantially increase worker productivity and improve performance in a variety of knowledge-intensive tasks ([Noy & Zhang, 2023](#); [Brynjolfsson et al., 2025](#)). At the same time, evidence indicates that the technology has been adopted rapidly across occupations and industries following the public release of ChatGPT ([Bick et al., 2024](#)). Despite this rapid diffusion, evidence from Denmark suggests that measurable effects on wages and employment have so far been limited ([Humlum](#)

& Vestergaard, 2025). Consequently, there remains considerable uncertainty regarding how generative AI affects labor market outcomes. This study contributes to this emerging literature by examining whether highly AI-exposed sectors in Sweden experienced different developments in wages and employment following the emergence of generative AI technologies.

This paper applies AI Industry Exposure (AIIE) scores to Swedish Standard Industrial Classification (SNI) sections in order to examine the potential effects of AI exposure on key sector-level labor market characteristics. These SNI sections represent broad industry groups within the Swedish economy. The exposure scores are converted from NAICS to the European Union's NACE classification system and subsequently to the Swedish SNI system. The study was initially designed to compare the sectors with the highest and lowest levels of AI exposure. However, due to data limitations, the analysis instead focuses on the sectors with the second highest and third lowest exposure scores. The development of these sectors key characteristics is then compared using a differences-in-differences (DiD) event study specification. The SNI sector with the third lowest AI exposure is identified as sector F, *Construction*, while the sector with the second highest exposure is identified as sector M, *Judicial, economic, scientific, and technical activities*. The analysis examines whether sectors with higher average AIIE scores experience different developments in hourly wages and employment levels, reflecting their greater exposure to AI adoption. This relationship is economically relevant because generative AI has the potential to alter the demand for labor, change the composition of workplace tasks, and affect worker productivity. Such changes may influence both wage formation and employment levels within affected sectors. Understanding whether highly exposed sectors develop differently from less exposed sectors therefore provides insight into how emerging AI technologies may reshape labor market outcomes and contribute to broader economic change.

Sweden could be a relevant context for studying the labor market effects of generative artificial intelligence due to its high degree of digitalization and comparatively strong labor market institutions. Sweden consistently ranks among the most digitally advanced economies in Europe, with high levels of internet usage, technological infrastructure, and digital adoption among firms and households. The Swedish labor market is also characterized by extensive collective bargaining coverage, relatively strong worker protections, and coordinated wage-setting institutions. Such institutions may influence how technological change affects labor market outcomes compared to economies with less coordinated markets, such as the United States. Strong institutions can potentially mitigate short-run labor market disruptions associated with generative

AI adoption, while Sweden's technologically advanced industries and service sectors may facilitate relatively rapid implementation of AI technologies. Consequently, the labor market effects of AI exposure observed in the Swedish context may differ from those identified in US-based studies, making Sweden an especially relevant case for examining the early economic implications of generative artificial intelligence ([OECD, 2018](#); [European Commission, 2022](#)).

2. Background

2.1 Generative AI

The generative AI systems implemented in workplaces are primarily based on large language models (LLM) and are commonly applied to tasks that are repetitive and language intensive, such as customer service, software development, content generation, and document processing. LLMs are software designed to process sequential data and predict the next word that follows in a sequence based on what is written before. The models are capable of this by training on large bodies of preexisting text, such as Wikipedia articles or digitized books. These AI tools can be further specialised to generate content for specific tasks and occupations. Research shows that these systems are particularly effective in tasks involving language production, information retrieval, summarization, and routine cognitive work ([Brynjolfsson et al., 2025](#)).

2.2 AI exposure scores

Felten, Raj, and Seamans construct an AI Occupational Exposure (AIOE) measure to estimate how exposed different occupations are to AI technologies, which will be necessary to calculate the AI Industry Exposure scores used in this study. The measure is based on ten selected AI applications listed by the Electronic Frontier Foundation (EFF), including image recognition, language modelling, translation, speech recognition, image generation, and reading comprehension. The EFF assigns each AI application a value between zero and one. Although these applications do not represent every possible use of AI, the authors argue that they capture many of the most important and common AI capabilities relevant to the labor market.

To connect AI technologies to workplace tasks, the selected AI applications are linked to all 52 occupational abilities defined in the US Occupational Information Network (O*NET) database. These occupational abilities include memorization, reaction time, stamina, and depth perception. The abilities are defined using survey responses collected from Amazon's Mechanical Turk survey, then evaluating how related each AI application is to each occupational ability. Thus, estimating a relatedness score between zero and one for every AI application–ability combination.

Felten et al. (2021) then calculate an ability-level AI exposure score by summing the relatedness scores across all ten AI applications. This approach assumes that

each AI application contributes equally to occupational exposure and that the applications affect abilities independently of one another, such that:

$$A_{ij} = \sum_{i=1}^{10} x_{ij}$$

Where A estimates the ability level exposure, which is indexed by the AI application i and the occupational ability j . A equals the sum of all application-ability relatedness scores, x , based on the Mechanical turk survey data.

Occupational exposure is finally calculated using occupational data from O*NET. The ability-level exposure scores are weighted according to how important and prevalent each ability is within a given occupation. Occupations that rely heavily on abilities strongly associated with AI applications therefore receive higher overall exposure scores, resulting in:

$$AIOE_k = \frac{\sum_{j=1}^{52} A_{ij} \times L_{jk} \times I_{jk}}{\sum_{j=1}^{52} L_{jk} \times I_{jk}}$$

In which AIOE is the AI occupational exposure score which is indexed by the occupation k . i indexes the AI application, j indexes the occupational ability and A_{ij} represents the ability level exposure calculated in the previous equation. The ability level scores were then weighted by the abilities prevalence, L_{jk} , and importance, I_{jk} , of each occupation measured in the O*NET database.

The industry-level measure of AI exposure (AIIE) used in this study is constructed by aggregating occupational AI exposure scores across occupations within each industry. Specifically, the AIIE measure is calculated as an employment-weighted average of occupational AI exposure scores, using 2019 employment data based on the four-digit North American Industry Classification System (NAICS) codes (Felten et al., 2021).

2.3 Industry classification systems

The North American Industry Classification System (NAICS) is a standardized industry classification system used in the United States, Mexico, and Canada to categorize businesses and economic activities into industries. The system assigns numerical industry codes to economic entities through a hierarchical structure in which classifications become more specific as additional digits are added to each code (Bureau of Labor Statistics, 2023). The Statistical Classification of Economic Activities (NACE) system is employed by member countries of the

European Union and functions as a classification tool for different economic activities. The system is structured very similarly to the North American system; however, the economic sectors are lettered from A to U rather than numbered. The classification thereby becomes more specific as additional digits are added to each sector code ([European Commission, n.d.](#)). Finally, the SNI system employed in Sweden also serves to classify economic activity within the country. SNI is based on the European NACE system which essentially makes the two interchangeable for economic analysis. This analysis will use SNI classifications published in 2007 rather than those published in 2025. The 2007 iterations were used because they acted as the official industry classification system in Sweden throughout most of the study period. This ensures consistent industry definitions over time and avoids potential measurement errors that could arise from applying the revised SNI 2025 classification retrospectively. Furthermore, the use of a stable classification framework improves comparability with existing literature and reduces the risk that observed effects are driven by industry reclassifications rather than underlying economic changes. This study is conducted at the broadest level of aggregation, namely the letter level sector level, which consists of 21 distinct economic sectors ([Statistics Sweden, n.d.](#)).

3. Methodology

3.1 Data

The study uses data gathered from Statistics Sweden (SCB). SCB is a Swedish government agency with the purpose of gathering and recording Swedish statistics, covering a broad range of economic, demographic, and social statistics. The data covers the time period of 1980 to 2025; however, some variables were observed at longer time periods than others. To avoid omitted observations due to missing values in the regression, the final sample is restricted to periods where all variables have complete data. The majority of the data is originally at a monthly frequency. To reduce short-term volatility and align the variables across datasets, the data are aggregated to quarterly observations for the regression estimations. The final dataset captures the time period of 2020 to 2025.

The regression is performed for two dependent variables. The average hourly wage in Swedish kronor for both SNI sectors of interest is the first dependent variable. The total employment level of both sectors is the second dependent variable. Both regressions include a total of six control variables, capturing ongoing short-term and long-term employment, each sector's gender distribution, and the GDP production approach of each sector.

Table 1 displays descriptive characteristics of both dependent variables, `hourly_wage` and `employed`, and the six control variables used in the study. The table shows all variables mean value, standard deviation, minimum and maximum value, and their observation count. The dataset has a total observation count of 432. The 144 observations shared between all variables are the ones used in the main analysis to avoid omissions in the regression specifications due to missing values.

Table 1. Descriptive statistics of dependent and control variables.

Variable	Mean	St. dev	Min	Max	Count
hourly_wage	174.1421	25.23811	126.4	237.1	432
employed	316582.1	31517.74	247060	372967	258
long_empl.	273441.2	16980.66	234572	300848	144
short_empl.	73105.78	7223.187	54897	93030	144
male_empl.	249879.1	62817.56	170431	332779	144
female_empl.	96667.94	51387.39	41895	161983	144
cap_formation	9394.208	2041.928	933	15664	432
prod_approach	76706.81	17237.82	38767	118503	432

The dataset created by [Felten et al. \(2021\)](#) is used for the purpose of determining which SNI sectors are the most and least exposed to artificial intelligence using their constructed formulas. The utilized dataset lists AI exposure scores for all occupations and industries listed by the US Bureau of Labor Statistics, together with their associated NAICS codes.

3.2 Industry classification conversion

To ensure that the measured AI industry exposure scores could be applied to Swedish industries, the first step of this study is to convert the American North American Industry Classification System (NAICS) codes into Swedish Standard Industrial Classification (SNI) codes. Because the SNI system is based on the European Statistical Classification of Economic Activities (NACE), the NAICS codes were first converted to NACE before being translated into SNI codes. It should be noted that no official conversion method exists between the NAICS and NACE classification systems. The conversion method employed in this study therefore relied on matching NAICS sectors to their corresponding NACE sectors based on the industries covered within each classification. Due to the structural similarities between the NAICS and NACE systems, all sectors contain highly similar industry compositions, resulting in a largely direct one-to-one

correspondence between the matched sectors. This makes the conversion process relatively straightforward and suitable for the purposes of this study.

The conversion from the NACE system to SNI is performed using the same approach. This conversion was similarly straightforward, as the Swedish SNI system is based directly on and coordinated with the EU's NACE classification system to ease economic analysis.

To conduct the analysis, the SNI sectors with the highest and lowest levels of AI exposure are next identified. These sectors are subsequently compared using a differences-in-differences regression to estimate the effects potentially associated with high AI exposure. To determine the relevant sectors, the average AI Industry Exposure (AIIE) score is calculated for the industries contained within each economic SNI sector using the industry level AIIE measures developed by Felten et al. (2021). The AIIE scores ranged from approximately -2 to 2, where higher values indicate greater exposure to artificial intelligence.

The economic sector with the lowest average exposure to artificial intelligence is identified as SNI sector A, *Agriculture, forestry and fishing*, with an average AIIE score of -1.558. The highest exposure is observed in SNI sector K, *Financial and insurance activities*, with an average score of 2.052. However, due to data limitations and other considerations discussed later, the analysis is instead conducted using the sectors with the second highest and third lowest exposure levels. The sector with the third lowest average AIIE score is SNI sector F, *Construction*, with a score of -0.997, while the second highest exposure is found in sector M, *Judicial, economic, scientific, and technical activities*, with an average score of 1.595.

3.3 Empirical strategy

The empirical strategy is based on a difference-in-differences (DiD) event study design and is performed in the statistics program Stata. The validity of this approach relies on the parallel trends assumption, which states that the treated and control sectors would have displayed similar outcome trends in the absence of treatment. Another assumption important to the DiD framework is the Stable Unit Treatment Value Assumption (SUTVA). Two requirements need to be fulfilled for SUTVA to apply. First, it requires that there are no spillover effects, meaning that the treatment one unit receives cannot affect any outcomes for a different unit. Second, SUTVA requires that there are no hidden variations of treatment, meaning that all treated units must be subjected to identical treatments. The DiD method compares the development of the sector average hourly wage and employment level between the highly exposed SNI sector M and the much less

exposed SNI sector F. The event study specification is estimated and illustrated separately for each outcome variable. Sector M, *Judicial, economic, scientific, and technical activities* will constitute the treatment group due to its greater exposure to AI adoption as implied by the sector containing a high average AIIE score. Sector F, Construction will then act as the control group as it is considerably less likely to experience AI adoption due to its lower average exposure score.

The post-treatment period of the study begins in November 2022, marked by the public release of ChatGPT, an LLM-based generative AI application developed by OpenAI. This period will span 13 quarters, from 2022q4 to 2025q4. The post-treatment begins here as ChatGPT becoming public marked a major breakthrough in the opportunities for widespread adoption and public awareness of generative AI technologies (Farzad et al., 2024).

3.4 Regression equations

This equation represents the Difference-in-Differences specification used to estimate the average treatment effect:

$$\begin{aligned} \text{hourlywage}_{it} &= \alpha + \beta_1 \text{DiD}_{it} + \beta_2 \text{treated}_i + \theta X_{it} + \delta_t + \varepsilon_{it} \\ \text{employment}_{it} &= \alpha + \beta_1 \text{DiD}_{it} + \beta_2 \text{treated}_i + \theta X_{it} + \delta_t + \varepsilon_{it} \end{aligned}$$

The dependent variables are *hourly_wage* and *employment*. While i indexes the two SNI sectors M, *Judicial, economic, scientific, and technical activities* and F, *Construction*. The current period is indexed by t , spanning observed 24 quarters. *DiD* is a binary interaction term that takes the value 1 for observations in the treated sector M from November 2022 onward and 0 otherwise. Its coefficient represents the estimated Difference-in-Differences treatment effect. *treated* is a dummy variable that takes the value 1 if it is observed for the treated sector M, and 0 otherwise. X_{it} is a vector of all control variables, capturing employment composition, gender distribution, capital formation, and production approach.

Quarterly fixed effects, δ_t , are included to help isolate the treatment effect by controlling for macroeconomic events and shocks that affect both SNI sectors equally over time. These include events such as inflation, nationwide technological changes, and economic recessions, while also accounting for many of the shocks associated with the COVID-19 pandemic. Unit fixed effects are excluded from the regression since the analysis only studies two SNI sectors. Heteroskedasticity-robust standard errors are used in the regression to account for potential heteroskedasticity in the error terms.

The event study segment is then performed within the DiD regression for the purpose of estimating how the treatment effect develops over time, to better evaluate the parallel trends assumption, and to observe whether treatment effects appear immediately or over time.

The event study model then appears as follows:

The dependent variables are *hourly_wage* and *employment*. While i indexes the two SNI sectors M, *Judicial, economic, scientific, and technical activities* and F, *Construction*. The current period is indexed by t , spanning 24 quarters. k indexes the event time relative to the treatment period, taking a value between -11 to 12 . *treated* is a dummy variable that takes the value 1 if it is observed for the treated sector M, and 0 otherwise. *eventquarter_t* is a series of dummy variables that represent each period relative to the treatment date and take the value 1 when an observation belongs to a specific quarter before or after treatment for the treated sector, and 0 otherwise. The period immediately before treatment ($k = -1$) is omitted to avoid perfect collinearity and serves as the reference category, meaning regression coefficients are estimated relative to this omitted pre-treatment time period. (*treated_i × eventquarter_t*) is a summation term including a series of interaction variables between the treated dummy and *event_quarter* dummies for each period relative to the omitted treatment period. β_k captures the effect of treatment at the event time k , relative to the omitted reference period. The set of interactions between treated and *event_quarter* are stored in a STATA macro named '*eventquarter*', which is used in the regression equation. θX_{it} is a vector of all control variables.

Quarterly fixed effects, δ_t , are included to help control for macroeconomic events and shocks that affect both SNI sectors equally over time. Such as inflation, nationwide technological changes, and economic recessions, while also accounting for many of the shocks associated with the COVID-19 pandemic. Unit fixed effects are excluded from the regression since the analysis only studies two SNI sectors. Heteroskedasticity-robust standard errors are used in the regression to account for potential heteroskedasticity in the error terms.

3.5 Control variables

The inclusion of long-term and short-term employment controls aims to account for differences in workforce composition across sectors. Sectors with varying proportions of long-term and short-term employees may exhibit different wage

and employment dynamics independent of AI exposure. Controlling for these factors helps isolate the estimated treatment effect from changes related to labor market attachment and employment structure.

Gender composition controls are included to account for potential differences in wage levels and employment patterns associated with the distribution of male and female workers across sectors. Since sectors differ in their gender composition and generative AI exposure may be unevenly distributed across occupations, controlling for gender composition reduces the risk that estimated treatment effects capture underlying demographic differences rather than the effects of AI exposure.

Capital formation is included as a control variable because investment activity may influence both wage and employment outcomes. Sectors experiencing higher levels of investment may simultaneously exhibit greater productivity growth, technological adoption, and labor market developments independent of generative AI exposure. Controlling for capital formation therefore helps separate the effects of AI exposure from broader investment-related developments.

A measure of productivity is included to account for differences in sectoral economic performance. More productive sectors may experience stronger wages and different employment trajectories regardless of their exposure to generative AI technologies. By controlling for productivity, the analysis aims to reduce omitted variable bias arising from underlying differences in sector performance and efficiency.

3.6 Graphical structure

The `coefplot` command is used to plot and graph the estimated coefficients for the event-time indicators for the 24 observed quarters relative to treatment. The event time indicators are arranged on the x-axis, while the estimated effect per period relative to the omitted time period ($t = -1$) is displayed on the y-axis. The coefficients represent the difference in the dependent variable between the treated and control sector in a specific period relative to the omitted reference period. The vertical lines attached to each coefficient represent 95 percent confidence intervals, displaying statistical uncertainties in the coefficient estimates. The event study plots constitute a visual test of the parallel trends assumption and also depict how the treatment effect develops over time.

4. Results

Figure 1 displays the event study estimates for hourly wage. The coefficients before treatment stray considerably away from zero and have imprecise confidence intervals. This suggests evidence of differing pre-treatment trends in hourly wage between the treated SNI sector M, and the control sector F. The treatment effect is estimated to be negative and appears to increase in magnitude over time. Meaning the treated sector M, is estimated to experience weaker wage developments relative to the control sector F, during the period following treatment. However, the estimated coefficients begin to decrease in value 3 quarters before treatment, implying that an effect separate from the treatment might partially cause a decline in wage developments.




Figure 1. DiD event study graph, hourly wage.

Figure 2 displays the event study estimates for employment. The estimated pre-treatment coefficients have large confidence intervals and tend to deviate from zero. Consequently, the parallel trends assumption receives limited empirical support. Furthermore, the wide confidence intervals and variation in the pre-treatment coefficients make it difficult to assess whether the treated and control sectors followed similar employment trends prior to treatment. Following treatment, the estimated coefficients are generally negative, suggesting relatively

weaker employment developments in sector M compared to sector F. However, this pattern is less pronounced than in the previous event study, and the large confidence intervals make it difficult to draw firm conclusions regarding the magnitude and persistence of the estimated treatment effect. Again, the coefficients are estimated to decrease in value 3 quarters before treatment, suggesting that a different effect causes relatively weaker developments in employment levels for the treated sector as well




Figure 2. DiD event study graph, employment level.

Table 2 displays the regression output with hourly wage as the dependent variable, with additional control variables being added in subsequent specifications to test robustness. The main coefficient, DiD, is negative in all estimations, significant at the 1 percent level throughout, and decreases in magnitude as more controls are introduced. The negative coefficient suggests that high exposure to artificial intelligence is associated with lower average hourly wages in the treated sector M, Judicial, economic, scientific, and technical activities. The reduction in the estimated effect following the inclusion of additional controls indicates that part of the wage differential is explained by workforce composition and sector-specific characteristics. The reduction in the estimated effect after adding more controls suggests that the earlier specifications may have suffered from omitted variable bias. Some of the initial wage differences appear to be explained by workforce composition and sector-specific

characteristics that were not included in the simpler models. However, the coefficient remains statistically significant across all specifications, indicating that the negative relationship is not entirely driven by these omitted factors.

Table 2. Regression table, hourly wage.

VARIABLE	(1) hourly_wage	(2) hourly_wage	(3) hourly_wage	(4) hourly_wage
DiD	-11.33*** (1.014)	-6.170*** (1.351)	-6.966*** (1.720)	-3.684*** (1.390)
treated	-17.78*** (0.548)	-26.55*** (1.513)	-36.55*** (12.61)	-45.09*** (9.187)
long_empl.		-0.000281*** (4.13e-05)	0.501 (0.331)	0.466 (0.299)
short_empl.		4.46e-05 (4.78e-05)	0.502 (0.331)	0.467 (0.299)
male_empl.			-0.502 (0.331)	-0.467 (0.299)
female_empl.			-0.502 (0.331)	-0.466 (0.299)
cap_formation				-0.00188*** (0.000328)
prod_approach				-0.00024*** (5.40e-05)
Constant	212.7*** (2.561)	289.9*** (13.67)	293.2*** (14.65)	314.6*** (13.49)
Time fixed effects	Yes	Yes	Yes	Yes
Observations	144	144	144	144
R-squared	0.972	0.982	0.983	0.987

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 3 presents the regression results with employment level as the dependent variable. Additional control variables are included in subsequent specifications to test the robustness of the estimates. The main coefficient, DiD, varies considerably across specifications and becomes less statistically significant as more controls are added. While the estimated coefficient is generally positive, the

substantial variation in its magnitude and the declining level of statistical significance make it difficult to draw firm conclusions about the effect of AI exposure on relative employment levels. These results suggest that the estimated treatment effect is sensitive given the current control variables used and should therefore be interpreted with caution.

Table 3. Regression table, employment level.

VARIABLE	(1) employed	(2) employed	(3) employed	(4) employed
DiD	22,637*** (2,256)	-1,814*** (312.7)	536.4** (229.5)	389.0* (226.5)
treated	(1,024)	(336.5)	(1,954)	(1,817)
long_employed	(1,024)	(336.5)	(1,954)	(1,817)
short_employed		0.960*** (0.00686)	48.77 (49.09)	40.00 (51.49)
male_employed		0.920*** (0.0126)	48.72 (49.09)	39.96 (51.49)
female_employed			-47.72 (49.09)	-38.95 (51.49)
cap_formation				-48.04 (49.09)
prod_approach				0.163*** (0.0608)
				0.00142 (0.0123)
Constant	362,918*** (4,802)	9,176*** (2,234)	-6,083*** (1,427)	-8,143*** (1,721)
Time fixed effects	Yes	Yes	Yes	Yes
Observations	144	144	144	144
R-squared	0.910	0.999	1.000	1.000

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

5. Discussion

5.1 Interpretations

The results suggest that Swedish economic sectors with high AI exposure were associated with relatively weaker wages following the introduction of generative artificial intelligence. However, the absence of parallel pre-treatment trends and indications of changes in labor market outcomes shortly before treatment weaken the causal interpretation of these findings. As a result, it is difficult to determine whether the observed changes were caused by the treatment itself. Furthermore, the SUTVA requirements are almost definitely not fully satisfied. Due to issues with the study design addressed previously, there are unavoidable spillover effects as both the control and treated sector receive varying degrees of treatment.

The employment results are considerably more uncertain. The estimated treatment coefficients were inconsistent across specifications, accompanied by large confidence intervals and generally low statistical significance. Similar to the wage analysis, indications of changes in employment prior to treatment further complicate causal interpretation. SUTVA would here also be violated because of weaknesses in the study design. Consequently, the results provide limited evidence that generative AI had a meaningful effect on employment levels during the study period.

The pre-treatment developments in labor market outcomes suggest that factors unrelated to generative AI may have influenced the results, raising concerns that the estimated treatment effect is overstated. Consequently, the causal impact of generative AI remains uncertain, particularly for employment, where no clear mechanisms can be inferred. Nevertheless, several AI-related mechanisms may explain the weaker wage developments observed in the treated sector. Generative AI can automate certain tasks and workflows, potentially reducing labor demand and strengthening employers' bargaining positions in wage negotiations. Additionally, firms may face short-term adjustment costs associated with AI implementation, while any productivity gains from the technology may require time to translate into higher wages.

There are several factors unrelated to AI adoption that may have influenced wage developments during the study period. First, the COVID-19 pandemic and the subsequent economic recovery likely affected the treated and control sectors differently. The pandemic altered labor demand, working conditions, and economic activity across industries, and these effects may have contributed to divergent wage trends prior to the introduction of generative AI. Broader

macroeconomic developments, including rising inflation, higher interest rates, and fluctuations in overall economic activity, may have influenced the sectors in different ways. These conditions likely affected firms' costs, investment decisions, and wage-setting processes. Third, sector-specific differences in labor demand may also help explain the observed patterns. The sectors differ in their sensitivity to business cycles and in the services they provide, meaning that changes in demand for professional services versus construction activity may have driven divergent developments in wages independent of AI adoption. Together, these factors provide plausible alternative explanations for the observed divergence in pre-treatment trends.

5.2 Results in relation to existing literature

The findings of this study align with recent empirical evidence on generative AI adoption. [Humlum and Vestergaard \(2025\)](#) find limited effects on wages and employment in Denmark despite rapid adoption of large language models. Similarly, [Brynjolfsson et al., \(2025\)](#) show that generative AI can improve worker productivity, particularly in communication and information-processing tasks. Together, these studies suggest that productivity gains may not immediately translate into measurable labor market effects. The absence of clear employment effects in the present study is therefore consistent with generative AI still being in an early stage of adoption.

The results can be contrasted with earlier research on automation and industrial robots. [Acemoglu and Restrepo \(2020\)](#) find that exposure to industrial robots reduced wages and employment in affected U.S. labor markets. While the negative wage estimates in this study are similar, the weak employment effects suggest that generative AI may influence labor markets differently. Unlike industrial robots, which mainly replace routine manual tasks, generative AI primarily affects cognitive and language-intensive work, implying different mechanisms and outcomes.

Additional evidence on generative AI adoption is provided by [Bick et al., \(2024\)](#), who document the rapid spread of generative AI tools following the public release of ChatGPT. Their findings indicate that AI was adopted across a wide range of occupations within a short period of time. This supports the use of the post-2022 period as a meaningful point at which AI exposure may have begun to influence labor market outcomes, even if measurable effects remain limited in the short run.

The results can also be interpreted through the framework of [Acemoglu and Restrepo \(2019\)](#), who argue that technological change both displaces existing

tasks and creates new ones. The overall effect on wages and employment depends on the balance between these forces. Applied to generative AI, this suggests that productivity-enhancing and task-creating mechanisms may offset reductions in labor demand, potentially explaining the weak employment effects observed in this study.

5.3 Limitations

The conversion from NAICS to NACE and subsequently to SNI classifications introduces potential measurement errors. The primary concern is that the AI exposure scores used in this study are calculated using the task compositions of occupations in the United States. Equivalent occupations within the European Union or Sweden may exhibit different task compositions, which could result in discrepancies in the estimated levels of AI exposure. Such differences would then affect the calculated AIIE scores and potentially alter which sectors are more exposed than others.

It also needs to be highlighted which and how many economic sectors this study analyses. As stated earlier, the study investigates the sectors that are second most exposed and third least exposed to AI. This is due to a combination of higher ranked sectors containing few industries or lacking sufficient variable data. The most and least exposed SNI sectors K, Financial and insurance activities and A, Agriculture, forestry, and fishing, lacked sufficient data from Statistics Sweden and were therefore opted out of the study. The second least exposed sector H, Transport and storage, was not chosen due to it containing few industries. Too few industries within a sector might give its average AIIE score lower statistical strength. This selection presents an under estimation of the treatment intensity in the analysis. The difference in average AI exposure scores between the most and least exposed sectors was approximately 3.6, ranging from -1.558 to 2.052 . However, the final sector selections shared a smaller exposure gap of approximately 2.6, ranging from -0.997 to 1.595 . Consequently, the treated and control sectors are relatively more similar in their estimated levels of AI exposure. A larger exposure differential between sectors would likely increase the contrast between the treatment and control sectors, potentially resulting in stronger estimated treatment effects and more pronounced deviations in the event study analysis.

Another limitation of the study is the small number of sectors included in the analysis. Since the estimation of the treatment effect relies on variation across a limited set of aggregated SNI sectors, the statistical power of the analysis is reduced, making the estimates potentially more sensitive to sector-specific

developments. These events that affect single sectors may disproportionately affect the regression results. Studying aggregated sectors risks overlooking heterogeneity within each sector. Industries and occupations within the same sector can differ considerably in their true exposure to AI technologies. The estimated treatment effect might not capture this variation in exposure at the industry or occupational level.

Critically, there is no definite separation between the treated and control group, causing a case of treatment contamination in the study. Because one SNI sector is expected to experience considerably more AI adoption due to its high exposure score, does not mean that the much less exposed sector experiences no AI adoption. The control group in this case is then also being treated, just presumably much less. An expected consequence of this is that the treatment effect likely becomes estimated as lower than it would if a truly untreated entity acted as the control group in the study.

Quarterly data was used instead of monthly observations in order to reduce the influence of short-term volatility and seasonal fluctuations in the labor market. Since the implementation of AI technologies is likely to occur gradually, quarterly observations may more effectively capture adjustments in wages and employment compared to higher frequency data. Quarterly data might however conceal more short-run effects that might potentially be identified with monthly observations. The choice to use quarterly data also drastically lowers the number of observations in the dataset. Had the study used monthly data instead, up to 72 observations would be observed rather than the now 24, which reduces statistical power and makes estimates less precise.

5.4 Methodological improvements

Several methodological improvements could have strengthened this study and potentially provided more robust evidence on the relationship AI exposure and labor market outcomes in Sweden. First, the analysis could have been conducted using less aggregated data. The broad coverage of SNI sectors risks obscuring potentially substantial variation between industries and occupations within the same sector. Utilizing data on the industry or occupational level would have likely produced more precise estimates of the effect following high AI exposure.

The study would have employed AI exposure measures based on Swedish occupational data, had they been readily available, rather than relying on exposure scores developed for the United States labor market. Because occupational task compositions likely differ somewhat between countries, using an exposure

measure based on Swedish occupational classifications could have reduced measurement error and improved the validity of the treatment variable.

Studying a larger number of treated and control sectors instead of focusing on only two sectors would have benefitted the study. Incorporating more sectors would have increased the statistical power of estimations while reducing the influence of sector-specific shocks unrelated to AI adoption.

Finally, using a direct measure of AI adoption for Swedish economic entities would improve the study. While exposure scores capture the potential for AI use, they do not necessarily reflect actual implementation. Data on AI investments, software expenditures, AI-related job postings, or firm-level surveys on AI adoption could have provided a more direct measure of AI use and strengthened the interpretation of the results.

6. Conclusion

This thesis examines the relationship between AI exposure and labor market outcomes across Swedish SNI sectors using a differences-in-differences event study design. The results suggest that highly AI-exposed sectors experienced relatively weaker wages following the emergence of generative AI technologies. In contrast, the analysis does not provide clear evidence that AI exposure has affected employment levels. Overall, the findings provide limited evidence that labor market outcomes in highly AI-exposed sectors have developed differently from those in less exposed sectors following the emergence of generative AI technologies.

The findings of the study should be interpreted with caution due to several methodological limitations. First, the post-treatment period remains relatively short, meaning the analysis likely captures only early-stage labor market effects of generative AI adoption. Furthermore, the study relies on only two aggregated SNI sectors, which may reduce statistical precision and conceal heterogeneity within industries and occupations. The conversion of AI exposure measures from NAICS classifications to SNI classifications may also introduce measurement error, as occupational task compositions may differ between the United States and Sweden. Additionally, the use of quarterly data may obscure short-term effects that could potentially be observed using higher-frequency data. Importantly, the estimated coefficients for both outcome variables begin to decrease in value shortly before the treatment is applied. This suggests that an effect separate from the treatment is affecting the relative development of employment levels and hourly wages in the treated group. Finally, the relatively small exposure difference between the treated and control sectors may reduce the contrast in AI adoption intensity, likely affecting the estimated treatment effects.

Despite the limitations of the study, the findings highlight the importance of continuing research on the labor market effects of generative artificial intelligence. As AI technologies become increasingly integrated into workplaces, understanding how different sectors and workers are affected will become more relevant for both policymakers and firms. This situation may develop differently in a context like Sweden, where high levels of digitalization and strong labor market institutions may influence how AI implementations impact wages and employment.

References

- Bick, Alexander. Blandin, Adam. Deming, David, J. (2024). *The Rapid Adoption of Generative AI*. National Bureau of Economic Research Working Paper No. 32966. <https://www.nber.org/papers/w32966>
- Autor, David, H. (2015). *Why Are There Still So Many Jobs? The History and Future of Workplace Automation*. Journal of Economic Perspectives. 29 (3), 3-30. <https://www.aeaweb.org/articles?id=10.1257/jep.29.3.3>
- Brynjolfsson, Erik. Li, Danielle. Raymond, Lindsey, R. (2025). *Generative AI at Work*. The Quarterly Journal of Economics. 140 (2), 889-942. <https://www.nber.org/papers/w31161>
- Eloundou, Tyna. Manning, Sam. Mishkin, Pamela. Rock, Daniel. (2023). *GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models*. arXiv preprint arXiv:2303.10130. <https://arxiv.org/abs/2303.10130>
- Felten, Edward. W. Raj, Manav. and Seamans, Robert. (2021) *Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses*. Strategic Management Journal. 42 (12), 2195-2217 <https://sms.onlinelibrary.wiley.com/doi/full/10.1002/smj.3286>
- Acemoglu, Daron. Restrepo, Pascual. (2020). *Robots and Jobs: Evidence from US Labor Markets*. The University of Chicago press Journals. 128 (6), 2188–2244. <https://www.journals.uchicago.edu/doi/full/10.1086/705716>
- Acemoglu, Daron. Restrepo, Pascual. (2019). *Automation and New Tasks: How Technology Displaces and Reinstates Labor*. Journal of Economic Perspectives, 33 (2), 3-30. <https://www.aeaweb.org/articles?id=10.1257/jep.33.2.3>
- Noy, Shakked. Zhang, Whitney. (2023) *Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence*, Science. 381 (6654), 187-192. <https://www.science.org/doi/10.1126/science.adh2586>
- Humlum, Anders. Vestergaard, Emilie. (2025). *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. National Bureau of Economic Research Working Paper. <https://www.nber.org/papers/w33777>
- OECD (2018), *OECD Reviews of Digital Transformation: Going Digital in Sweden*, OECD Reviews of Digital Transformation, OECD Publishing, Paris. <https://doi.org/10.1787/9789264302259-en>
- European Commission (2022) *Digital Economy and Society Index (DESI) 2022 Sweden*. Sweden in the Digital Economy and Society Index. <https://digital-strategy.ec.europa.eu/en/policies/desi-sweden>
- U.S. Bureau of Labor Statistics (n.d.) *North American Industry Classification System (NAICS) at BLS*. <https://www.bls.gov/bls/naics.htm>. Accessed 12 March 2026.
- European Commission (n.d.) *NACE – Statistical classification of economic activities*. Eurostat. Available at: <https://ec.europa.eu/eurostat/web/nace/overview>. Accessed 12 March 2026.

- Statistics Sweden (SCB) (n.d.) *Standard for Swedish Industrial Classification (SNI)*.
<https://www.scb.se/dokumentation/klassifikationer-och-standarder/standard-for-svensk-naringsgrensindelning-sni/>. Accessed 11 March 2026.
- Kamrani, Farzad. Kanestad, Linus. Limér, Christoffer. Pelzer, Björn. Smedberg, Iza. Tegen, Agnes. Wickenberg Bolin, Ulrika. (2024). *Large Language Models in Defense: Challenges and Opportunities*. Totalförsvarets forskningsinstitut (FOI). FOI-R--5544--SE. <https://www.foi.se/rapportsammanfattning?reportNo=FOI-R--5544--SE>

Publishing and archiving

Approved students' theses at SLU can be published online. As a student you own the copyright to your work and in such cases, you need to approve the publication. In connection with your approval of publication, SLU will process your personal data (name) to make the work searchable on the internet. You can revoke your consent at any time by contacting the library.

Even if you choose not to publish the work or if you revoke your approval, the thesis will be archived digitally according to archive legislation.

You will find links to SLU's publication agreement and SLU's processing of personal data and your rights on this page:

- <https://libanswers.slu.se/en/faq/228318>

☒ YES, I, Jesper Lindbom, have read and agree to the agreement for publication and the personal data processing that takes place in connection with this

☐ NO, I/we do not give my/our permission to publish the full text of this work. However, the work will be uploaded for archiving and the metadata and summary will be visible and searchable.