WORKING PAPER 10/2024 (ECONOMICS)

Artificial Intelligence, Hiring and Employment: Job Postings Evidence from Sweden

Erik Engberg, Mark Hellsten, Farrukh Javed, Magnus Lodefalk, Radka Sabolová, Sarah Schroeder and Aili Tang

ISSN 1403-0586

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

Artificial Intelligence, Hiring and Employment: Job Postings Evidence from Sweden*

Erik Engberg   Mark Hellsten   Farrukh Javed   Magnus Lodefalk†
Radka Sabolová   Sarah Schroeder   Aili Tang

November 7, 2024

Abstract

This paper investigates the impact of artificial intelligence (AI) on hiring and employment, using the universe of job postings published by the Swedish Public Employment Service from 2014-2022 and universal register data for Sweden. We construct a detailed measure of AI exposure according to occupational content and find that establishments exposed to AI are more likely to hire AI workers. Survey data further indicate that AI exposure aligns with greater use of AI services. Importantly, rather than displacing non-AI workers, AI exposure is positively associated with increased hiring for both AI and non-AI roles. In the absence of substantial productivity gains that might account for this increase, we interpret the positive link between AI exposure and non-AI hiring as evidence that establishments are using AI to augment existing roles and expand task capabilities, rather than to replace non-AI workers.

Keywords: Artificial Intelligence; Technological Change; Automation; Labour Demand.

JEL Codes: D22, J23, J24, O33.

*We thank David Autor, Dan Johansson, Ina Jäkel, Hildegunn Kyvik-Nordås, Martin Längkvist and Manav Raj for feedback and Edvin Ahlander and Lukas Ma for research assistance. Engberg, Hellsten, Lodefalk and Sabolová acknowledge financial support from the Ratio Institute, Lodefalk from the Jan Wallanders och Tom Hedelius Foundation (grant P19-0234) and Torsten Söderberg Foundation (grant E46/21), and Sabolová from the Jean Monet Network and the Erasmus+ Program of the European Union.

†Corresponding author: Magnus Lodefalk, Associate Professor. Address: Department of Economics, Örebro University, SE-70182 Örebro, Sweden, Telephone: +46 19 303407, +46 722 217340; Global Labor Organization, Essen, Germany; Ratio Institute, Stockholm, Sweden. E-mail: *magnus.lodefalk@oru.se*.

1. Introduction

We investigate the impact of artificial intelligence (AI) on recent changes in establishments' hiring patterns, exploiting the universe of job postings from the Swedish Public Employment Service (*Arbetsförmedlingen*, AF) and universal register data for Swedish firms. Our study is motivated by three main facts: (i) recent breakthroughs in AI technologies, such as natural language processing, have enabled machines to perform or assist with tasks traditionally carried out by white-collar workers, including roles in paralegal and consulting work (Agrawal *et al.*, 2018, Zhang *et al.*, 2021, Dell'Acqua *et al.*, 2023); (ii) labour demand for AI-related skills has substantially increased over last decade (Maslej *et al.*, 2024); and (iii) widespread societal and academic concerns persist regarding AI's potential negative impact on labor markets, particularly on white collar employment (Frey and Osborne, 2017, Korinek and Stiglitz, 2017, Acemoglu and Johnson, 2023, Susskind, 2022, Susskind and Susskind, 2018).

AI has the potential to redefine the boundary between codified versus tacit knowledge in the workplace. By automating cognitive, nonroutine tasks like making predictions, AI can take over specific functions while augmenting human roles in areas that are less amenable to automation, such as product innovation and customer engagement. When automation predominates without significant productivity gains, worker displacement may occur; conversely, if productivity gains are substantial, or AI complements human work or enables the creation of new tasks, it may lead to increased labour demand (e.g., Acemoglu *et al.*, 2022). Ultimately, the question of the employment implications of AI is an empirical one.

While research on digital automation and labour markets has primarily studied the use of computers or robots in manufacturing, our study also encompasses service firms. Advances in AI increasingly expose professionals and other service workers to 'thinking machines' that can perform complex, cognitive tasks. Moreover, the services sector now employs a growing share of the workforce, and, in many economies, professional service industries comprise a larger portion of employment than manufacturing.[1]

Our study also contributes to the emerging literature that addresses the scarcity of data on AI adoption in labor markets by using vacancy data, as pioneered in the U.S. context (Zolas *et al.*, 2021). The study most closely related ours is Acemoglu *et al.* (2022), which utilises online vacancy data (2010-2018) to estimate the impact of AI exposure on establishment hiring in US industries that use AI. They find a positive (negative) impact on the hiring of workers with AI (non-AI) skills, but no employment effects. Building on this seminal work, we construct, to the best of our knowledge, the first non-US establishment dataset (2014-2022) that links AI exposure to hiring at the establishment level by leveraging Swedish job postings and universal register data. Sweden is a small, open, highly servicified and digitalised economy with similar AI adoption rates as in the USA. The Swedish register data enable us to measure AI exposure based on the actual workforce composition of establishments. We further supplement our analysis with detailed survey data on AI use from Statistics Sweden.

We document a sharp rise in demand for AI-related skills and a strong, positive and statistically significant association between AI exposure and AI hiring. In contrast to Acemoglu *et al.* (2022), we also find that the more AI-exposed establishments increase their non-AI hiring, resulting in overall employment growth. The positive link between AI exposure and non-AI hiring in Sweden would be consistent with AI complementing workers in tasks, new tasks being introduced, and/or the substantial productivity effects from AI as found in recent studies (e.g., Hirvonen *et al.*, 2022, Acemoglu and Restrepo, 2019, Acemoglu *et al.*, 2022, Dell'Acqua *et al.*, 2023). Exploring potential mechanisms, we cautiously conclude that Swedish establishments may be using AI to augment, rather than replace, non-AI workers.

[1]Such services, which, e.g., include legal, auditing, management, architectural, and advertising services, are also increasingly important in manufacturing as it servicifies, i.e., increasingly use, produce and sell services (Lodefalk, 2013, 2017, Arnarson and Gullstrand, 2022)

2. Data and empirical framework

2.1. Data Description

We employ the job vacancies posted on Sweden's largest recruitment site *Platsbanken.* The site is run by AF and has a mean of 448,000 jobs posted each year. Posting job ads at AF was mandatory, but, since 2008, only remained so for central government establishments. However, the regulatory change has only slightly reduced postings at AF (Cronert, 2019). Posting job ads is free, and ads may be reposted to other sites, e.g., LinkedIn. From each job ad, we use detailed information: job title, occupational code, organisation, municipality, and specific skill requirements. Skill requirements are used to establish whether a vacancy is AI-related.

We consider an AI-related vacancy as one requiring at least one AI skill. We extend the categorisation of keywords by AI skill used by Deming and Noray (2020) by merging it with keywords from Alekseeva *et al.* (2020) and the OECD work of Baruffaldi *et al.* (2020) (see Table A1). As displayed in Figure 1, the share of AI vacancies has increased almost exponentially in both Sweden and the USA since the mid-2010s, although from low levels.

For our study, we aggregate the job vacancy data to the establishment level, pooling vacancies by organisation and municipality. We then split the data into two time periods to reduce noise and improve precision: 2014-2016 and 2019-2022. The first period captures the state right before major AI breakthroughs, e.g., Google's notable improvement in machine translation when adopting deep neural networks late 2016, and the second captures the state thereafter, while allowing organisations time to act upon these advances. Finally, we study changes in the posting of vacancies (establishment hiring) and employment. Table A.1 in the Online Appendix provides further descriptive statistics at the establishment level.

We relate establishment changes in AI hiring and non-AI hiring to initial exposure to AI. The exposure variable, also used in Acemoglu *et al.* (2022), is based on the AI occupational

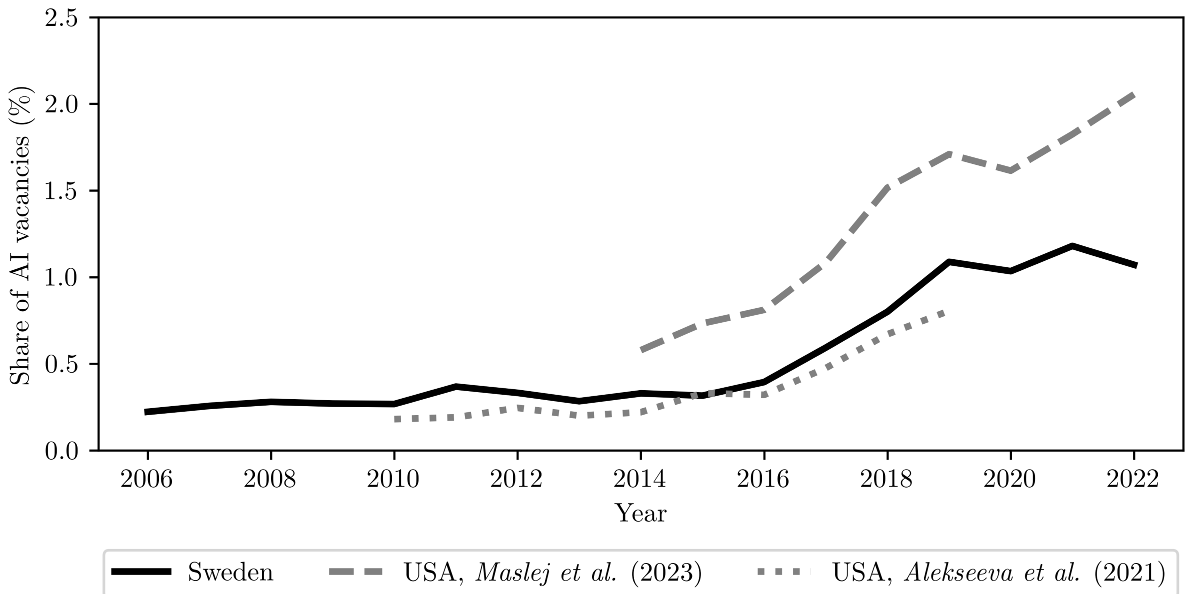

FIGURE 1
Share of AI Vacancies

Notes: This figure displays the shares of AI vacancies in the AF data from 2006 and onwards, as well as the timelines for the US data from Maslej *et al.* (2023) and Alekseeva *et al.* (2021). Potential differences in methodology are not controlled for.

exposure index from (Felten *et al.*, 2018), which assigns a score to each detailed occupation, representing the likelihood that the occupation is affected by recent AI advancements. Exposure is measured via the ability of AI to perform the work of an occupation. To construct the exposure variable, the index for each establishment is calculated as the weighted average of the index values for all occupations employed in the establishment between 2014 and 2016, based on the universal Longitudinal Integrated Database for Health Insurance and Labour Market Studies (LISA) from Statistics Sweden.[2] While the exposure measure is positively correlated with the occupational and establishment shares of AI vacancies in Sweden, the measure is *ex ante* agnostic about its impacts: whether more exposed establishments will adopt AI, how they will do it (by hiring or externally sourcing AI-services), and why (e.g., to replace or augment non-AI workers).

[2]LISA includes all individuals (≥ 15 years old) living in Sweden.

2.2. *Empirical Estimation*

To study the impact of AI on establishment hiring and employment, we build on Acemoglu *et al.* (2022) and estimate the following regression model:

$$\Delta Y_{i,t_1-t_0} = \beta_0 + \beta_1 AI_{i,t_0} + \mathbf{X}_{i,t0}\boldsymbol{\beta}_{\boldsymbol{X}} + \epsilon_{i,t_1-t_0} \tag{1}$$

where $\Delta Y_{i,t_1-t_0}$ is the change in the inverse hyperbolic sine of outcome Y for establishment i between periods t_1 and t_0, AI_{i,t_0} is AI exposure, $X_{i,t0}$ is a row vector of confounders, the βs are regression parameters, and ϵ_{i,t_1-t_0} is an i.i.d. error term. For comparison, the AI exposure variable is standardised so that the regression parameter is interpreted as the change in the outcome variable associated with a one-standard deviation increase in the explanatory variable.

In essence, the specification relates changes in hiring or employment to initial conditions in terms of establishment workforce composition and the resulting exposure to AI developments. We then add on indicator variables to control for confounding factors related to establishment size, municipality, and firm. A potential concern is that AI exposure could be confounded by a positive correlation between exposure to AI and to other computer software. We therefore also add the measure of software exposure from Webb (2020), which is built on occupational task and patent data.

3. Results

3.1. *Impact of AI Exposure on Employment and Hiring*

In Table 1, we present our estimation of Equation (1) for the hiring of AI and non-AI workers as well as overall employment. Starting from a basic specification with covariates size and location, we find a positive and statistically significant association between AI exposure and both the hiring of AI and non-AI workers (Col. 1 and Col. 3) as well as employment (Col. 5).

We then control for location and firm heterogeneity, as well as for software exposure (Col's 2, 4, and 6). In this within-firm specification, which is our preferred one, the estimated links to AI-hiring and employment are larger, while the one to non-AI hiring is somewhat smaller. A one-standard deviation increase in AI exposure is linked to a 27 (23) percent increase in the hiring of AI (non-AI) workers, and a 5 percent increase in employment.

TABLE 1

AI Exposure and Changes in Vacancies/Employment

Dependent variable	ΔAI-hiring		ΔNon-AI-hiring		ΔEmployment	
	(1)	(2)	(3)	(4)	(5)	(6)
AI exposure	15.454***	26.561***	29.646***	23.153***	3.014**	4.507**
	(2.135)	(6.639)	(4.875)	(6.854)	(1.480)	(2.239)
Size FE	✓		✓		✓	
Location FE	✓	✓	✓	✓	✓	✓
Firm FE		✓		✓		✓
Observations	72,756	22,921	72,756	22,921	47,575	16,475

Notes: This table presents estimates from six establishment-level regressions, weighted by baseline employees. The outcome variable is the change ($\times 100$) in the inverse hyperbolic sine of AI vacancies, non-AI vacancies, and total employment. The AI exposure measure from Felten *et al.* (2018) is based on baseline employees, standardized. Columns (2), (4), and (6) include the software exposure covariate from Webb (2020). Observations are lower in specifications with firm fixed effects due to omitted singleton establishments. Standard errors are clustered at the firm level. * p<0.1; ** p<0.05; *** p<0.01.

Our results on AI hiring align qualitatively with those of Acemoglu *et al.* (2022) for the USA. However, while they find indications of a negative impact of AI exposure on non-AI hiring and an insignificant impact on aggregate employment, we find evidence of a positive impact on both non-AI hiring and overall employment at the establishment level. Their interpretation suggests that, in US establishments, AI displaces non-AI work. In contrast, our results imply that Swedish firms may be using AI to augment rather than replace non-AI workers–a point we revisit later. Another possible explanation is that Sweden's higher wages and labour scarcity may drive AI-induced productivity gains that support continued demand for labour (Acemoglu and Restrepo, 2019). To reconcile these differences, we note that while Sweden and the USA have similar AI adoption rates (5.4% and 6.6%, respectively), other structural differences between the two economies are substantial (SCB, 2020, Zolas *et al.*, 2021).

Starting with firm size, the average firm in Sweden has four employees, while the US firm has four times as many, 16 employees (Bisnode, 2024, U.S. Census Bureau, 2024). Moreover, the US job postings data in Acemoglu *et al.* (2022) are from the web scraping of firm websites by the company Lightcast (previously Burning Glass Technologies). Lightcast data are known to overrepresent high-skilled occupations and larger firms (Cammeraat and Squicciarini, 2021). Our job postings data are from the official, and previously mandatory, outlet for job postings in Sweden, why our data are likely to be less skewed towards high-skilled occupations and larger firms.

These differences in the data generating processes may at least partially explain the patterns we observe. Behaviour varies between larger and smaller firms in ways that may result in larger AI-exposed firms predominantly substituting AI workers for non-AI workers – in effect, specialising in AI – while smaller ones add on AI to complement existing non-AI workers. Firstly, for a large firm that adopts new technology, it is arguably easier to fire employees, than for a small firm where existing employees are necessary for daily operations.[3] Secondly, smaller firms are commonly family-owned, and may therefore behave differently when adopting AI than larger ones, which most often are publicly listed.[4]

3.2. Effects by Establishment Size

To examine potential heterogeneous effects across different size categories, we rerun our estimations of Equation (1) for establishments above and below the median size. In Table 2, we display the results.

The estimates in Panel A do not indicate that larger Swedish establishments decrease non-AI hiring, as in the US sample. Instead, the results are similar to the ones for all Swedish establishments. However, interestingly, in Panel B, we find that for the smaller establish-

[3]Smaller firms are less likely to shed labour in response to shocks, and, if young, they tend to grow faster (Bjuggren, 2015, Coad and Karlsson, 2022).

[4]In Sweden, 88 percent of micro-firms, which have at most 9 employees, are family owned, whereas only 8 percent of large firms are (Andersson *et al.*, 2018). Family ownership has been found to be associated with slower but more steady growth, being geographically dispersed, and being more likely to retain workers in times of crisis (Andersson *et al.*, 2018, Baù *et al.*, 2024).

Table 2

AI Exposure and Changes in Vacancies/Employment, by Establishment Size

	Panel A: > median establishment size					
Dependent variable	ΔAI-hiring		ΔNon-AI-hiring		ΔEmployment	
	(1)	(2)	(3)	(4)	(5)	(6)
AI exposure	16.408***	27.780***	30.730***	23.116***	3.569**	4.595*
	(2.282)	(7.120)	(5.248)	(7.313)	(1.570)	(2.366)
Size FE	✓		✓		✓	
Location FE	✓	✓	✓	✓	✓	✓
Firm FE		✓		✓		✓
Observations	36,863	15,553	36,863	15,553	26,581	12,114

	Panel B: < median establishment size					
Dependent variable	ΔAI-hiring		ΔNon-AI-hiring		ΔEmployment	
	(1)	(2)	(3)	(4)	(5)	(6)
AI exposure	0.293**	0.259	2.960***	11.090***	0.148	5.839**
	(0.144)	(0.645)	(0.715)	(3.848)	(0.474)	(2.796)
Size FE	✓		✓		✓	
Location FE	✓	✓	✓	✓	✓	✓
Firm FE		✓		✓		✓
Observations	35,893	5,776	35,893	5,776	20,994	3,550

Notes: This table displays estimates from twelve establishment-level regressions, with baseline establishment number of employees as weights. Throughout, the outcome variable is the change in the inverse hyperbolic sine of AI vacancies, non-AI vacancies, and number of employees, multiplied by 100. The regressor is the AI exposure measure of Felten *et al.* (2018), average of baseline establishment employees, normalised by its standard deviation. Estimations are performed on two different samples: establishments above median (8) baseline number of employees (Panel A), and below median baseline number of employees (Panel B). There are two regressions for each dependent variable. In Col's (2), (4) and (6), the software exposure measure of Webb (2020) is included as a covariate. Lower number of observations in specifications including firm fixed effects are due to singleton establishments being omitted. Lower number of observations in specifications including firm fixed effects in Panel B are due to smaller firms more often being singleton establishments. Standard errors are clustered at the firm level. * p<0.1; ** p<0.05; *** p<0.01.

ments, exposure to AI is no longer statistically significantly linked to AI hiring, while it still is for non-AI hiring and employment. This could mean that smaller establishments do not respond to AI exposure by adopting AI technology.[5] Alternatively, smaller establishments do adopt AI, but source their AI services from external suppliers (SCB, 2020). We explore this possibility , using recent stratified firm survey data from SCB (2020, 2023).[6] In Figure 2, we present stylised results from a probability model, regressing the probability of a firm using internal or external AI services on AI exposure.

[5] AI is primarily adopted by larger firms and firms who have already invested in, e.g., cloud computing, and smaller firms may lack the expertise for adopting AI (Alekseeva *et al.*, 2021, Zolas *et al.*, 2021, SCB, 2023).

[6] The survey ('ICT usage in enterprises') was mandatory and distributed to all firms with $\geq$ 200 employees and a stratified random sample of firms with $\geq$ 10 employees. The response rate was 88% and 83% , in 2019 and 2021, respectively, resulting in approximately 4,200 firms being included each year.

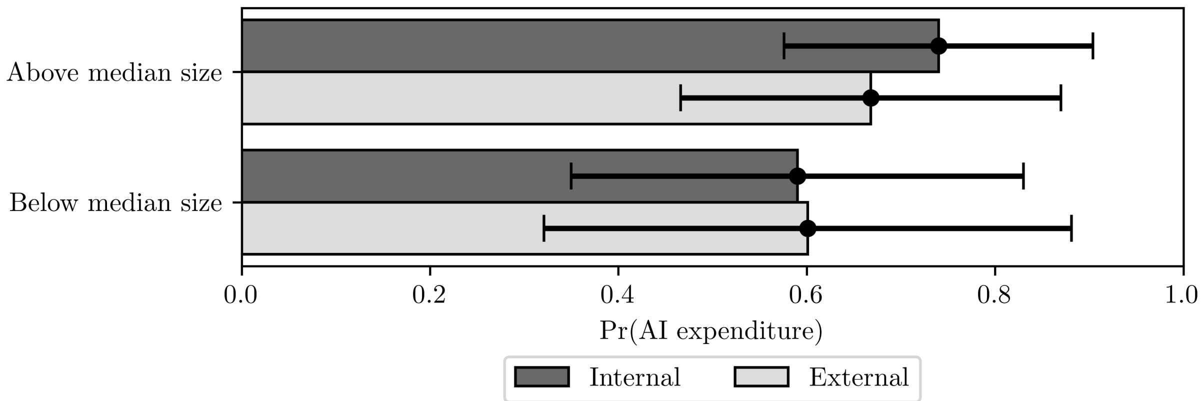

FIGURE 2

Exposure to AI and AI Use

Notes: This figure displays estimated coefficients from four firm-level probit regression. Throughout, the outcome variable is an indicator variable for using internally or externally sourced AI, defined as any AI expenditure in 2019 or 2021, using firm survey data from Statistics Sweden. The regressor is the AI exposure measure of Felten *et al.* (2018), based on firm employment in 2019, normalised by its standard deviation. Estimations are performed separately for firms below and above the mediannumber of employees. The sample is further limited to firms represented in the main regressions, in Table 1. Firm size fixed effects, 3-digit industry fixed effects, and the software exposure measure of Webb (2020) are included in all regressions. Error bars display the 95% confidence intervals.

For both larger and smaller firms, exposure to AI is positively and statistically significantly associated with the probability to use AI services.[7] We therefore conclude that AI exposure is associated with AI use also for small firms in Sweden. However, larger firms are more inclined to source AI services internally, a trend that does not hold for smaller firms, which rely more on external providers.

3.3. Alternative Mechanisms

Having ruled out firm size as the primary explanation for the differing results between the USA and Sweden, we turn to labour market characteristics. Sweden has a compressed wage structure and high *de facto* minimum wages, while the USA has a very dispersed wage distribution and low minimum wages.[8] As in many other OECD countries, in Sweden,

[7]Firms that spend on AI services are substantially more likely to hire AI workers, see Table A2. The link between spending on external AI services and hiring non-AI workers is also statistically significant, while only a fraction of the one for hiring AI workers. However, spending on internal AI services is not associated with the hiring of non-AI workers.

[8]The Gini coefficient is 30(40) for Sweden (the USA), and the population share with an income or consumption below 50% of the median is 11(16) for Sweden (the USA), based on income after taxes and benefits or consumption (World Bank, 2024). In 2023, for Sweden, the p10 to median wage was 73, and, for the USA, the minimum to mean (median) wage was 18 (26) (OECD, 2024a, SCB, 2024).

the labour market has been increasingly tight since the great financial crisis (Nordin and Hammarlund, 2024, OECD, 2024b). It is therefore possible that the high Swedish wages, in combination with labour scarcity, have led to substantial productivity improvements with AI-induced automation, thereby sustaining demand for non-AI workers.

To investigate this further, we re-estimate Equation (1), using total factor productivity as the outcome variable. The results are displayed in Col's (1)-(2) of Table A3. Interestingly, we find only a weakly significant and small positive association between AI exposure and productivity, and no significant association when controlling for industry heterogeneity. In the absence of a substantial productivity effect, we interpret the positive and statistically significant link between AI exposure and non-AI hiring as an indication that Swedish establishments are using AI to augment, rather than replace, non-AI workers, such as for developing new products or services.[9] This would also be in line with evidence from the previously mentioned survey that suggests that Swedish firms mainly use AI, for example, to develop new offerings, customer insights, and gain market shares, rather than to improve internal processes (SCB, 2020).[10] The stylised results in Col's (3)-(4) of Table A3 are consistent with this conjecture, showing a positive, albeit weakly significant, association between AI exposure and subsequent net turnover.

4. Concluding remarks

The importance of granular and multicountry evidence on AI and hiring patterns can hardly be overstated (Arntz *et al.*, 2017, Frank *et al.*, 2019, Zolas *et al.*, 2021). We exploit rich Swedish public job posting data and universal register data for services and manufacturing industries to investigate the establishment-level impact of AI on hiring and employment. We

[9]In Finland, Hirvonen *et al.* (2022) also find that advanced technologies adoption increases employment and the technologies are for providing new products.

[10]Most (46-60%) of the surveyed firms use AI to develop or improve offerings or gain insights on relations with customers/users. Only a minority (39%) state that AI use is related to internal processes or other purposes, with small firms mentioning it less frequently (31%) than large firms (57%). Not even for the large firms is the improvement of internal processes the main purpose of using AI.

find a positive association between exposure to AI and hiring of AI workers. Notably, in contrast to findings by Acemoglu *et al.* (2022) for the USA, we find significant and positive effects of AI exposure on both non-AI hiring and total employment in Sweden.

The absence of non-AI worker displacement or major productivity gains from AI adoption suggests that Swedish establishments are predominantly using AI to augment workers in their tasks. However, further research is needed to confirm these patterns and examine their persistence over time. Overall, we conclude that recent breakthroughs in AI technologies and the resulting surges in demand for AI skills do not appear to have significantly negatively impacted the employment of Swedish establishments, at least for now.

References

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). 'Artificial Intelligence and Jobs: Evidence from Online Vacancies.' *Journal of Labor Economics*, 40(S1), S293-S340.

Acemoglu, D., and Johnson, S. (2023). *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity.* Basic Books UK.

Acemoglu, D., and Restrepo, P. (2019). 'Automation and New Tasks: How Technology Displaces and Reinstates Labor.' *Journal of Economic Perspectives*, 33(2), 3-30.

Ackerberg, D., Caves, K., and Frazer, G. (2015). 'Identification Properties of Recent Production Function Estimators.' *Econometrica*, 83(6), 2411–2451.

Agrawal, A., Gans, J., and Goldfarb, A. (2018). *Prediction Machines: The Simple Economics of Artificial Intelligence.* Harvard Business Press.

Alekseeva, L., Giné, M., Samila, S., and Taska, B. (2020) 'AI Adoption and Firm Performance: Management versus IT.' https://ssrn.com/abstract=3677237

Alekseeva, L., Azar, J., Giné, M., Samila, S., and Taska, B. (2021). 'The demand for AI skills in the labor market.' *Labour Economics*, 71.

Andersson, F.W., Johansson, D., Karlsson, J., Lodefalk, M., and Poldahl, A. (2018). 'The characteristics of family firms: exploiting information on ownership, kinship, and governance using total population data.' *Small Business Economics*, 51, 539–556.

Arnarson, B. T., and Gullstrand, J. (2022). 'Linking local services to global manufactures.' *The Scandinavian Journal of Economics*, 124.

Arntz, M., Gregory, T., and Zierahn, U. (2017). 'Revisiting the risk of automation.' *Economics Letters*, 159, 157-160.

Baruffaldi, S., van Beuzekom, B., Dernis, H., Harhoff, D., Rao, N., Rosenfeld, D., Squicciarini, M. (2020), *Identifying and measuring developments in artificial intelligence: Making the impossible possible,* OECD STI WP 2020/05.

Baù, M., Karlsson, J., Haag, K., Pittino, D., and Chirico, F. (2024). 'Employee layoffs in times of crisis: do family firms differ?' *Entrepreneurship and Regional Development*, 36, 5-6, 722-744.

Bisnode (2024). Serrano Database. Retrieved from https://www.hhs.se/en/houseoffinance/data-center/serrano/.

Bjuggren, C.M. (2015). 'Sensitivity to shocks and implicit employment protection in family firms.' *Journal of Economic Behavior and Organization*, 119, 18-31.

Cammeraat, E., and Squicciarini, M. (2021). 'Burning Glass Technologies' data use in policy-relevant analysis: An occupation-level assessment.' OECD STI Working Papers, 2021/05.

Coad, A., and Karlsson, J. (2022). 'A field guide for gazelle hunters: Small, old firms are unlikely to become high-growth firms.' *Journal of Business Venturing Insights*, 17.

Cronert, A. (2019). 'Is regulatory compliance by employers possible without enforcement? Evidence from the Swedish labor market', IFAU WP 2019:23.

Dell'Acqua, F., McFowland III, E., Mollick, E., Lifshitz-Assaf, H., Kellogg, K.C., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K.R. (2023), 'Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality', Harvard Business School WP 24-013.

Deming, D., and Noray, K. (2020). 'Earnings Dynamics, Changing Job Skills, and STEM Careers.' *The Quarterly Journal of Economics*, 135(4), 1965-2005.

Felten, E., Raj, M., and Seamans, R. (2018). 'A method to link advances in artificial intelligence to occupational abilities.' *AEA Papers and Proceedings*, 108, 54-57.

Frank, M. R., Autor, D., Bessen, J. E., Brynjolfsson, E., Cebrian, M., Deming, D. J., ... and Rahwan, I. (2019). 'Toward understanding the impact of artificial intelligence on labor.' *Proceedings of the National Academy of Sciences*, 116(14), 6531-6539.

Frey, C. B., and Osborne, M. A. (2017). 'The future of employment: How susceptible are jobs to computerisation?' *Technological Forecasting & Social Change*, 114(2017), 254-280.

Hirvonen, J., Stenhammar, A., and Tuhkuri, J. (2022). 'New Evidence on the Effect of Technology on Employment and Skill Demand.' Unpublished manuscript, MIT.

Korinek, A., and Stiglitz, J. E. (2017). 'Artificial intelligence and its implications for income distribution and unemployment.' NBER WP 24174.

Levinsohn, J., and Petrin, A. (2003) 'Estimating Production Functions Using Inputs to Control for Unobservables.' *The Review of Economic Studies*. 70(2) 317–341

Lodefalk, M. (2013). 'Servicification of manufacturing–evidence from Sweden.' *International Journal of Economics and Business Research*, 6(1), 87-113.

Lodefalk, M. (2017). 'Servicification of firms and trade policy implications.' *World Trade Review*, 16(1), 59-83.

Maslej, N., Fattorini, L., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Ngo, H., Niebles, J., Parli, V., Shoham, Y., Wald, R., Clark, J., and Perrault, R.

(2023) 'The AI Index 2023 Annual Report.' AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, Stanford, CA, April

Maslej, N., Fattorini, L., Perrault, R., Parli, V., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J., Shoham, Y., Wald, R., and Clark, J. (2024). 'The AI Index 2024 Annual Report.' AI Index Steering Committee, Institute for Human-Centered AI, Stanford University, April.

Nordin, M., and Hammarlund, C. (2024). 'Arbetskraftsbrist: ett problem eller en möjlighet?' Report, AgriFood economics centre.

OECD (2024a). Earnings: Minimum wages relative to median wages. OECD Employment and Labour Market Statistics (database).

OECD (2024b). 'OECD Employment Outlook 2024: The Net-Zero Transition and the Labour Market.' Report, OECD Publishing, Paris,

SCB (2020). 'Artificial intelligence (AI) in Sweden 2019.' Report, Statistics Sweden.

SCB (2023). 'AI-use in the private and public sector.' Report, Statistics Sweden.

SCB (2024). Average salary and salary dispersion by sector, occupation (SSYK 2012) and sex, year 2023. Statistical database, Statistics Sweden.

Susskind, D., and Susskind, R. (2018). 'The Future of the Professions.' *Proceedings of the American Philosophical Society*, 162(2), 125-138.

Susskind, D. (2022). 'Technological Unemployment in Bullock.', forthcoming in Bullock, Justin, et al. (2022), *Oxford Handbook of AI Governance.*

U.S. Census Bureau. (2024). County Business Patterns CB2000CBP. Retrieved from https://data.census.gov/

Webb, M. (2020). 'The Impact of Artificial Intelligence on the Labor Market.' Unpublished manuscript, Stanford.

World Bank (2024). Poverty and Inequality Platform (version 20240326_2017 and 20240326_2011). World Bank Group, Data set.

Zhang, D., Mishra, S., Brynjolfsson, E., Etchemendy, J., Ganguli, D., Grosz, B., Lyons, T., Manyika, J., Niebles, J. C., Sellitto, M., Shoham, Y., Clark, J., and Perrault, R. (2021). 'The AI Index 2021 Annual Report.' AI Index Steering Committee, Human-Centered AI Institute, Stanford University, March.

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., ... and Dinlersoz, E. (2021). 'Advanced Technologies Adoption and Use by US Firms: Evidence from the Annual Business Survey.' NBER WP 28290.

Appendix

Table A1
AI Keywords

Deming and Noray (2020)			
Splunk	Support Vector Machines (SVM)	Unsupervised LearningCaffe Deep Learning Framework	Natural Language Processing
Apache Hadoop	Bayesian Networks	Boosting (Machine Learning)	Natural Language Toolkit (NLTK)
Sqoop	Clustering	Semi-Supervised Learning	Speech Recognition
Apache Hive	Cluster Analysis	Chef Infrastructure Automation	Pattern Recognition
MapReduce	Neural Networks	Automation Tools	Kernel Methods
TensorFlow	Convolutional Neural Network (CNN)	Automated Testing	Image Recognition
Scikit-learn	Recurrent Neural Network (RNN)	Automation Systems	Object Recognition
Mahout	Human Machine Interface (HMI)	Office Automation	Image Processing
Keras	Human Machine Interface (HMI) Control Systems	Automation Consulting	Machine Translation
OpenCV	Supervised Learning (Machine Learning)	Sales Automation Software	Text Mining
Xgboost	Machine-To-Machine (M2M) Communications	Automation Test Environment	Recommender Systems
Libsvm	Machine Code	Marketing Automation	Latent Semantic Analysis
Word2vec	Machine Vision	Laboratory Automation	Sentiment Analysis / Opinion Mining
Artificial Intelligence	Computer Vision	Automation Techniques	Virtual Agents
Machine Learning	Machine Translation (MT)	Automated Underwriting System	Chatbot
Robotics	Torch (Machine Learning)	Gradient Boosting	AI Chatbot
Decision Trees	Deep Learning	Random Forest	
Alekseeva *et al.* (2020)			
AI ChatBot	Image Recognition	MARF	Sentiment Analysis
AI KIBIT	IPSoft Amelia	MoSes	Opinion Mining
ANTLR	Ithink	MXNet	Sentiment Classification
Apertium	Keras	Natural Language Processing	Speech Recognition
Artificial Intelligence	Latent Dirichlet Allocation	Natural Language Toolkit	Supervised Learning
Automatic Speech Recognition	Latent Semantic Analysis	NLTK	Support Vector Machines
ASR	Lexalytics	ND4J	SVM
Caffe Deep Learning Framework	Lexical Acquisition	Nearest Neighbor Algorithm	TensorFlow
Chatbot	Lexical Semantics	Neural Networks	Text Mining
Computational Linguistics	Libsvm	Object Recognition	Text to Speech
Computer Vision	Machine Learning	Object Tracking	TTS
Decision Trees	Machine Translation	OpenCV	Tokenization
Deep Learning	MT	OpenNLP	Torch
Deeplearning4j	Machine Vision	Pattern Recognition	Unsupervised Learning
Distinguo	Madlib	Pybrain	Virtual Agents
Google Cloud Machine Learning Platform	Mahout	Random Forests	Vowpal
Gradient boosting	Microsoft Cognitive Toolkit	Recommender Systems	Wabbit
H2O	MLPACK	Semantic Driven Subtractive Clustering Method	Word2Vec
IBM Watson	Mlpy	SDSCM	Xgboost
Image Processing	Modular Audio Recognition Framework	Semi-Supervised Learning	
Baruffaldi *et al.* (2020)			
action recognition	face recognition	learning automata	robot
human action recognition	facial expression recognition	link prediction	biped robot
activity recognition	factorisation machine	logitboost	humanoid robot
human activity recognition	feature engineering	long short term memory (LSTM)	human-robot interaction
adaboost	feature extraction	lpboost	industrial robot
adaptive boosting	feature learning	machine intelligence	legged robot
adversarial network	feature selection	machine learning	quadruped robot
generative adversarial network	firefly algorithm	extreme machine learning	service robot
ambient intelligence	fuzzy c	machine translation	social robot
ant colony	fuzzy environment	machine vision	wheeled mobile robot
ant colony optimisation	fuzzy logic	madaboost	rough set
artificial intelligence	fuzzy number	MapReduce	rule learning
human aware artificial intelligence	fuzzy set	Markovian	rule-based learning
association rule	intuitionistic fuzzy set	hidden Markov model	self-organising map
autoencoder	fuzzy system	memetic algorithm	self-organising structure
autonomic computing	t s fuzzy system	meta learning	semantic web
autonomous vehicle	Takagi-Sugeno fuzzy systems	motion planning	semi-supervised learning
autonomous weapon	gaussian mixture model	multi task learning	sensor fusion
backpropagation	gaussian process	multi-agent system	sensor data fusion
Bayesian learning	genetic algorithm	multi-label classification	multi-sensor fusion
bayesian network	genetic programming	multi-layer perceptron	sentiment analysis
bee colony	gesture recognition	multinomial naive Bayes	similarity learning
artificial bee colony algorithm	gradient boosting	multi-objective optimisation	simultaneous localisation mapping
blind signal separation	gradient tree boosting	naive Bayes classifier	single-linkage clustering
bootstrap aggregation	graphical model	natural gradient	sparse representation
brain computer interface	gravitational search algorithm	natural language generation	spectral clustering
brownboost	hebbian learning	natural language processing	speech recognition
chatbot	hierarchical clustering	natural language understanding	speech to text
classification tree	high-dimensional data	nearest neighbour algorithm	stacked generalisation
cluster analysis	high-dimensional feature	neural network	stochastic gradient
cognitive automation	high-dimensional input	artificial neural network	supervised learning
cognitive computing	high-dimensional model	convolutional neural network	support vector regression
cognitive insight system	high-dimensional space	deep convolutional neural network	swarm intelligence
cognitive modelling	high-dimensional system	deep neural network	swarm optimisation
collaborative filtering	image classification	recurrent neural network	particle swarm optimisation
collision avoidance	image processing	neural turing	temporal difference learning
community detection	image recognition	neural turing machine	text mining
computational intelligence	image retrieval	neuromorphic computing	text to speech
computational pathology	image segmentation	non negative matrix factorisation	topic model
computer vision	independent component analysis	object detection	totalboost
cyber physical system	inductive monitoring	object recognition	trajectory planning
data mining	instance-based learning	obstacle avoidance	trajectory tracking
decision tree	intelligence augmentation	pattern recognition	transfer learning
deep belief network	intelligent agent	pedestrian detection	trust region policy optimisation
deep learning	intelligent software agent	policy gradient methods	unmanned aerial vehicle
dictionary learning	intelligent classifier	Q-learning	unsupervised learning
dimensionality reduction	intelligent geometric computing	random field	variational inference
dynamic time warping	intelligent infrastructure	random forest	vector machine
emotion recognition	Kernel learning	rankboost	support vector machine
ensemble learning	K-means	recommender system	virtual assistant
evolutionary algorithm	latent dirichlet allocation	regression tree	visual servoing
differential evolution algorithm	latent semantic analysis	reinforcement learning	xgboost
multi-objective evolutionary algorithm	latent variable	relational learning	
evolutionary computation	layered control system	statistical relational learning	

Table A2

AI Use and Hiring

Dependent variable	Pr(AI-hiring)		Pr(non-AI-hiring)	
Source of AI services	Internal	External	Internal	External
	(1)	(2)	(3)	(4)
AI use (1,0)	0.921***	0.901***	0.094	0.278**
	(0.115)	(0.160)	(0.101)	(0.140)
Controls	✓	✓	✓	✓
Observations	3,801	3,801	6,082	6,082

Notes: This table displays estimates from four firm-level probit regressions. Throughout, the outcome variables are the probabilities to post at least one AI vacancy or non-AI vacancy. The regressor is a dummy variable representing the use of internally or externally sourced AI, defined as any AI expenditure, and using survey data from Statistics Sweden. Controls (in logs) are human and capital intensities as well as labour productivity. The software exposure measure of Webb (2020) is also included as a covariate is all regressions. Standard errors are clustered at the firm-level. * p<0.1; ** p<0.05; *** p<0.01.

Table A3

AI Exposure, Productivity and Revenues

Dependent variable	ΔTotal factor productivity		ΔRevenue	
	(1)	(2)	(3)	(4)
AI exposure	4.939*	2.653	5.750	13.620*
	(2.918)	(2.468)	(7.539)	(7.546)
Size FE	✓	✓	✓	✓
Industry FE		✓		✓
Observations	20,680	20,665	57,647	57,636

Notes: This table displays estimates from four firm-level regressions, with baseline firm number of employees as weights. Throughout, the outcome variable is the change in the inverse hyperbolic sine of total factor productivity, and net revenue, multiplied by 100. The regressor is the AI exposure measure of Felten *et al.* (2018), average of baseline firm employees, normalised by its standard deviation. There are two regressions for each dependent variable. In Col's (2) and (4), the software exposure measure of Webb (2020) is included as a covariate. Regression is at firm level. Following the method of Table 1 as closely as possible, t_0 is 2014-2016, while t_1 is 2020 (2020 is the latest year of firm financial data). Total factor productivity is estimated using the methodology of Levinsohn and Petrin (2003), with corrections by Ackerberg *et al.* (2015). Lower number of observations in Col's (1)-(2) are due to the exclusion of observations with zeroes in the input variables for the estimation of total factor productivity, as these are log transformed. Standard errors are clustered at firm level. * p<0.1; ** p<0.05; *** p<0.01.

Online Appendix

Artificial Intelligence, Employment and Skills

Erik Engberg

Mark Hellsten

Farrukh Javed

Magnus Lodefalk

Radka Sabolová

Sarah Schroeder

Aili Tang

November 7, 2024

A. Additional Tables and Figures

Table A.1: Descriptive Statistics

	Obs.	Mean	Median	Std.
Establishment-level sample:				
Number of employees	89,445	38.171	7.667	314.563
Felten *et al.* (2018) AI exposure	89,433	0.357	< mean	0.028
Webb (2020) AI exposure	89,433	0.407	< mean	0.137
ΔAI-hiring	245,948	1.992	0	26.927
ΔNon-AI-hiring	245,948	16.005	0	216.431
ΔEmployment	61,982	8.280	< mean	80.905
Firm-level sample:				
Number of employees (2019)	69,997	54.068	6	573.456
Pr(Internal AI expenditure)	5,907	0.052	0	0.222
Pr(External AI expenditure)	5,907	0.025	0	0.156
Felten *et al.* (2018) AI exposure (2019)	69,966	-0.043	< mean	0.893
Webb (2020) AI exposure (2019)	69,966	-0.269	< mean	0.614
Number of employees (2014-2016)	76,405	52.412	6.333	547.663
ΔTotal factor productivity	20,712	-10.637	> mean	51.646
ΔNet revenue	63,316	-29.857	> mean	332.588
Felten *et al.* (2018) AI exposure (2014-2016)	76,394	-0.160	< mean	0.843
Webb (2020) AI exposure (2014-2016)	76,394	-0.127	< mean	0.662

Notes: This table displays summary descriptives of data at establishment and firm level. Medians consisting of micro-data replaced with size relative to mean.

Table A.2: Exposure to AI and AI use (extensive margin)

Dependent variable	Pr(AI-internal)			Pr(AI-external)		
Sample split by median	Full sample	Below	Above	Full sample	Below	Above
	(1)	(2)	(3)	(4)	(5)	(6)
AI exposure	0.677***	0.590***	0.740***	0.646***	0.601***	0.668***
	(0.068)	(0.120)	(0.082)	(0.086)	(0.140)	(0.101)
Size FE	✓			✓		
Industry FE	✓	✓	✓	✓	✓	✓
Observations	5,828	2,458	2,910	5,708	2,150	2,656

Notes: This table displays estimates for six probit regression specifications. Throughout, the outcome variable is probabilities to use internally or externally sourced AI, defined as any AI expenditure in the period 2019-2022. The regressor is the AI exposure measure of Felten *et al.* (2018), based on firm employees in 2019, normalised by its standard deviation. Estimations are performed on three different samples: The full sample of firms, firms below median (36) number of employees in 2019, and firms above median number of employees in 2019. The sample is further limited to firms present in the main regression in Table 1. The software exposure measure of Webb (2020) is included as a covariate is all regressions. Lower number of observations in the below median sample is lower due to more industries having no AI usage, and so they are omitted for perfectly predicting AI usage. Robust standard errors are in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

Table A.3: Exposure to AI and AI use (intensive margin)

Dependent variable	log(AI-internal expenditure)			log(AI-external expenditure)		
Sample split by median	Full sample	Below	Above	Full sample	Below	Above
	(1)	(2)	(3)	(4)	(5)	(6)
AI exposure	0.377***	0.183***	0.596***	0.197***	0.077***	0.333***
	(0.035)	(0.030)	(0.069)	(0.028)	(0.020)	(0.056)
Size FE	✓			✓		
Industry FE	✓	✓	✓	✓	✓	✓
Observations	5,897	2,941	2,956	5,897	2,941	2,956

Notes: This table displays estimates for six regression specifications. Throughout, the outcome variable is the log of expenditure on internally or externally sourced AI, using the sum of AI expenditure in the period 2019-2022. Log approximated by inverse hyperbolic sine to allow for zeroes. The regressor is the AI exposure measure of Felten *et al.* (2018), based on firm employees in 2019, normalised by its standard deviation. Estimations are performed on three different samples: The full sample of firms, firms below median (36) number of employees in 2019, and firms above median number of employees in 2019. The sample is further limited to firms present in the main regression in Table 1. The software exposure measure of Webb (2020) is included as a covariate is all regressions. Lower number of observations in the below median sample is lower due to more industries having no AI usage, and so they are omitted for perfectly predicting AI expenditure. Robust standard errors are in parentheses. * p<0.1; ** p<0.05; *** p<0.01.

Table A.4: AI Exposure and Financial Indicators

Dependent variable	ΔRevenue		ΔProfit		ΔEBITDA	
	(1)	(2)	(3)	(4)	(5)	(6)
AI exposure	5.750	13.620*	-59.951	-11.416	-14.816	116.619
	(7.539)	(7.546)	(98.907)	(82.276)	(96.432)	(72.666)
Size FE	✓	✓	✓	✓	✓	✓
Industry FE		✓		✓		✓
Observations	57,647	57,636	57,647	57,636	57,647	57,636

Notes: This table displays estimates from four firm-level regressions, with baseline firm number of employees as weights. Throughout, the outcome variable is the change in the inverse hyperbolic sine of net revenue, profit, and Earnings before interest, taxes, depreciation and amortization (EBITDA), multiplied by 100. The regressor is the AI exposure measure of Felten *et al.* (2018), average of baseline firm employees, normalised by its standard deviation. There are two regressions for each dependent variable. In Col's (2), (4) and (6), the software exposure measure of Webb (2020) is included as a covariate. Regression is at firm level. Following the method of Table 1 as closely as possible, t_0 is 2014-2016, while t_1 is 2020 (2020 is the latest year of firm financial data). Standard errors are clustered at firm level. * p<0.1; ** p<0.05; *** p<0.01.

References

Felten, E., Raj, M., and Seamans, R. (2018). 'A method to link advances in artificial intelligence to occupational abilities.' *AEA Papers and Proceedings*, 108, 54-57.

Webb, M. (2020). 'The Impact of Artificial Intelligence on the Labor Market.' Unpublished manuscript, Stanford.