[NOT FOR DISSEMINATION]
What is an AI worker? Evidence from Matched Vacancy-Hires*

Mark Hellsten
mark.hellsten@uni-tuebingen.de

June 1, 2026
Link to latest version

Abstract

Despite growing policy interest in AI workforce development, a key open question is whether AI workers constitute a distinct set of occupations or a self-selected subset within them. Using a novel matched vacancy-hire dataset of Swedish job postings and employer-employee register data, I study the characteristics and wages of workers who develop and apply AI. Observable traits including high cognitive ability, tertiary degree in STEM subjects, and prior AI experience largely explain sorting into AI-related occupations, but do little to distinguish these workers from their occupational peers. I then use earnings as a window onto unobservable differences. AI workers earn approximately 5% more than their peers, a premium explained entirely by firm-level wage premia, suggesting unobservable traits make them attractive to high-paying employers. AI workers with doctorate degrees stand out among their occupational peers. Taken together, the supply of average AI workers can be controlled via STEM education participation. The elite within-occupation AI talent are largely created via researchers moving to private sector.

Keywords: Artificial Intelligence; Wage premium; Labour demand; Skills; Technology diffusion

JEL Codes: D22, J24, M2, O33.

*Paper previously circulated as 'Returns to AI skills'. I am grateful to seminar participants at the Ratio Institute, Aarhus University, University of Chicago Booth School of Business, and University of Chicago, conference participants at Labor Market Inequality and Workplace AI 2025 (Aarhus University), AI and the Future of Work 2024 (AI at Wharton), Internal FIND Workshop 2024, Aarhus-Kiel Workshop 2023, and DGPE workshop 2023, as well as Mathias Barding, Rosario Crinò, Georg Graez, Anna Gumpert, Anders Humlum, Ina Jäkel, Michael Koch, Magnus Lodefalk, Sarah Schroeder, Bastian Schulz, Lars Skipper, and Emilie Vestergaard for valuable comments and discussions. Financial support from the European Union through the European Research Council (ERC-2022-SG ORGANDICT, No. 101076790), the Carlsberg Foundation, and the Ratio Institute is gratefully acknowledged.

1 Introduction

To benefit from the productivity increases associated with AI technologies, firms need workers with the skills to use or develop the technology. Despite reported scarcities (OECD, 2023; SCB, 2023a), as well as increased policy focus on AI training across multiple countries[1], not much is known about what AI workers are.[2] Workers with AI skills might range from simply a group of occupations working with AI (AI occupations), to a within-occupation group of technology pioneers, responsible for developing and enabling the uptake of the technology. This question has important policy implications related to how an economy can supply AI workers, as well as whether these workers pose a key barrier to the adoption of AI.

In a method of direct matching of skill requirements in job vacancies to individuals hired to fill the vacancies, I employ job vacancy data from Sweden, along with register data on the entire Swedish population. This is one of the first papers employing this approach, which allows for several key contributions not possible when relying solely on one of these data sources. Mainly, tying AI skills to individuals allows for the analysis of employment history, educational background, as well as character traits of these workers. This enables a analysis of what constitutes an AI worker both in terms of observable traits, and unobservables via wages. This allows me to shed light on whether these workers primarily consist of AI occupations, or a within-occupation group of AI pioneers, with important implications on

[1] OECD (2024) lists public training initiatives for the development of AI professionals being launched in Austria, Croatia, Germany, Korea, Norway, Poland, Singapore, United Kingdom, and the United states (With additional countries limited to training in general AI literacy.). In Sweden the government has adopted an AI strategy plan, which includes labor skill initiatives such as the funding and monitoring of university courses on the topic of AI for professionals, initiatives to increase the number of individuals choosing STEM fields for their education, and mandates for regional assessments of skill demand (Government Offices of Sweden, 2026).

[2] In the context of this paper, AI skills encompass the specialized abilities required for the development and application of Artificial Intelligence. Specifically, skills exclusive to the development and application of AI are considered, such as the ability to comprehend and create an image recognition algorithm. Skills used for other tasks, like general programming, are not included in this definition, neither is the use of AI-based tools in cases where knowledge of the workings of the tool are not required, such as use of Generative AI chatbots. I refer to workers with these skills as AI hires and AI workers.

the state of scarcity of AI workers, and how AI workers are created.[3]

I first employ the observable characteristics of educational background, cognitive ability, prior AI experience, as well as migrant background to investigate whether these traits are related to workers with AI skills, and how much of this can be explained by workers sorting into AI-related occupations. In terms of education, AI occupations tend to consist of tertiary educated individuals with a science or ICT degree. Workers in AI occupations are also related to individuals with high cognitive ability, with prior AI experience, as well as higher likelihood of being a recent immigrant. However, these traits explain sorting into AI occupations, rather than higher levels compared to their occupational peers. AI workers only stand out beyond their occupational peers in terms of higher levels of education, science degree, and prior AI experience.

Using worker earnings as a proxy to understand the role of unobservable traits among AI workers, I find that AI workers earn 5% more than their occupational peers, entirely explained by employment in high paying firms. AI workers with a doctorate degree have a much higher premium, 20% compared to doctorates without AI skills, and 9-12% more than AI workers with tertiary education.

Taken together, the average AI worker is mostly defined as workers in occupations with tertiary education, in sciences and ICT fields. However, the supply of a more elite and highly demanded within-occupation group of AI workers comes from AI workers with a researcher background. This implies that the policymaker can increase the supply of workers in AI occupations via increasing the number of individuals in STEM educations, to letting domestic AI researchers spill over to private sector, the latter enabling more advanced AI development.

This paper is related to the small but growing literature on AI skills and AI workers (Alek-

[3]Additional advantages include the possibility to control for characteristics of the hired individuals, mitigating potential biases in the earnings stemming from candidate expectations not explicitly stated in the job vacancies, as well as matching filtering out job vacancies that were not successfully filled, thus limiting the sample to actual employment.

seeva, Azar, Giné, Samila, & Taska, 2021; Firpo, Niemann, & Danilov, 2025; Green & Lamby, 2025; Pouliakas, Santangelo, & Dupire, 2025; Squicciarini & Nachtigall, 2021). This literature provides foundational descriptives of job vacancies requiring AI skills, including the occupational and sectoral distribution, with most but not all documenting increasing demand for AI skills and a wage premium. This paper contributes with information on who the workers possessing these skills are, whether they are more defined by their occupation or within-occupation skills, as well as how these workers can be supplied to firms in need of AI skills. These papers on AI skills fit within the broader literature on the role of pioneering workers in the diffusion of new technologies (Babina, Fedyk, He, & Hodson, 2023; Bloom, Hassan, Kalyani, Lerner, & Tahoun, 2021; Goldfarb, Taska, & Teodoridis, 2023; Tambe, 2014).

In addition, this paper is related to the literature on the effects of AI on the labor market. These papers often implicitly or explicitly includes the effect of AI on AI workers. Acemoglu, Autor, Hazell, and Restrepo (2022); Bonfiglioli, Crinò, Gancia, and Papadakis (2024); Engberg et al. (2025); Hellsten, Khanna, Lodefalk, and Yakymovych (2025) finds positive connections between AI and demand for AI workers.[4]

Lastly, the paper ties into the task-based framework for the effects of automation on employment and earnings in relation to tasks and skills (Acemoglu & Autor, 2011; Acemoglu & Restrepo, 2018, 2019). These papers propose a theoretical framework wherein technological advancements render certain low-skilled work tasks obsolete while simultaneously creating new tasks. The recent emergence and subsequent surge in the demand for AI skills[5], suggests that AI skills are utilized to complete work tasks created by new technologies. This paper contributes with empirical evidence on the nature of the workers performing newly emerging complex work tasks.

[4]See also (Webb, 2020) for effects on high skilled jobs. AI workers would make up a complementary subsample of these workers. See also (Brynjolfsson, Mitchell, & Rock, 2018; Felten, Raj, & Seamans, 2018) for evidence on labor market effects of AI not specifically highlighting certain workers.

[5]As seen in Figure **??** in the Appendix

The paper is outlined as follows: In Section 2 I discuss the data used, including the methodology of matching job vacancies to individuals. Section 3 provides details on who AI workers are, including an analysis of to what degree observable characteristics are related to AI occupation and a within-occupation group of AI workers. Section 4 analyzes unobservable traits and the returns to AI skills.I conclude with Section 5.

2 Data

The main methodology of this paper consists of identifying vacancies requiring AI skills and matching them with register data on the individuals hired to fill the vacancies. This section details the extraction of skills from job advertisements, the sample and characteristics of the register data, and the matching process.

2.1 AI skills and job vacancies

In the context of this paper, AI skills refer to the skills needed specifically to develop or use AI.

In order to access skills required for positions, I utilize online job vacancy data sourced from platsbanken, the online vacancy platform of the Swedish Public Employment Service. This is among the largest (if not the largest) job vacancy platforms in Sweden, hosting an average of one million job vacancies annually, when including single advertisements seeking to fill multiple vacancies. To contextualize this number, the Swedish labor force comprised approximately 5.5 million individuals in 2021 (SCB, 2021). The extensive use of this platform is partly attributed to a past regulatory requirement for firms to post vacancies via the Swedish Public Employment Service. Although this regulation was phased out in 2007, the platform did not experience a significant decline in the number of job postings thereafter (Cronert, 2019). From this data, I extract skill content mentioned in the texts of the vacancies, including AI skills.

To identify job vacancies that necessitate AI skills, I employ a comprehensive list of AI skills compiled from three distinct sources (Alekseeva, Gine, Samila, & Taska, 2020; Baruffaldi et al., 2020; Deming & Noray, 2020).[6,7] These sources provide a wide array of skills ranging from general concepts like 'machine-learning' to more specialized platforms such as 'tensorflow' (an open-source machine learning library), and specific machine learning applications like 'action recognition' (which involves using machine learning to identify human movements and actions).

To ensure that only parts of job vacancy texts describing skill requirements are used in the classification, I use a Keras deep learning Tensorflow model via the JobAd Enrichments API provided by the Swedish public employment office.[8] The embedding layer is trained on roughly six million job vacancy texts with the purpose of being able to identify words that, within the context of the text, describe skills requirements. This avoids misclassifications where AI related words are interpreted as skill requirements when mentioned in unrelated sections of the vacancy, e.g. the company description. The result is that each job vacancy is assigned an accompanying list of identified competencies, along with a probability of how likely it is that the term is requested by the employer. I keep competencies with a probability > 0.5. I then use text matching against the list of AI keywords to create an indicator variable representing whether a job vacancy requires AI skills. Figure 1 displays the most common individual AI competencies demanded in job vacancies.

2.2 Register data and sample

My study also incorporates register data that encompasses the full population of individuals and firms in Sweden. The register data provides background information on hired individuals,

[6]However, I exclude the skill 'Python' from Deming and Noray (2020), despite its prevalence. Python is a versatile programming language with broad applications beyond AI, and therefore, its inclusion would not accurately represent a pure AI skill.

[7]The AI keywords in this paper predates major breakthroughs in generative AI. Hence, proficiency in using tools powered by AI are not considered AI skills as defined in this paper.

[8]This API is used in the search engine of platsbanken to provide relevant vacancies for search terms.

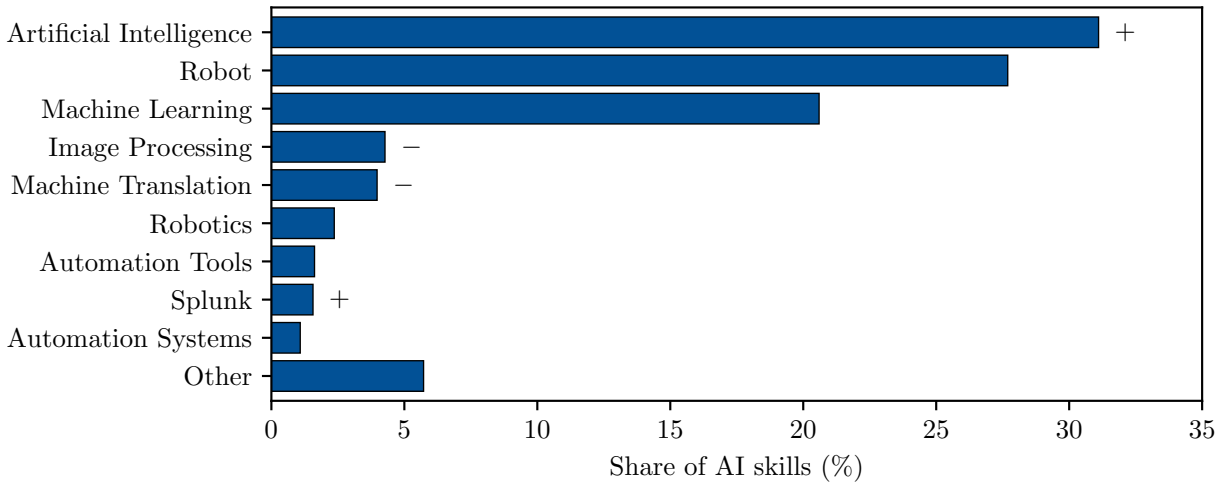


FIGURE 1
Top AI skills

Notes: This figure displays the 9 most common Alekseeva et al. (2020); Baruffaldi et al. (2020); Deming and Noray (2020) AI skill keywords from vacancies posted in 2023. Plus (minus) signs denotes the AI skills that are among the 10 skills with the highest growth (negative growth) between 2014-2023.

including earnings, work history, educational background, and more.[9]

The sample of matched data spans from 2014 to 2023. As the data is being matched to job vacancies, I only include new external hires. Thus, incumbent workers as well as internal hires are not considered. Some sample restrictions are imposed with two primary objectives: First, I make restrictions chosen to give the analysis of AI workers a representative control group, while also balancing the sample. This includes limiting the sample to persons in the ages 19 to 65, firms with minimum five employees, and to hires in industries that had at least one AI hire in the full time period. Second, as the earnings is an important outcome variable in the empirical analysis of this paper, I exclude individuals who do not have a representative value of full-time annual earnings. This involves excluding workers who were not employed in the hiring firm for at least one full calendar year after their employment. As there is no way to identify part-time workers, these workers would also have unrepresentative earnings, and so they are not desirable to include in the analysis. In an attempt to filter out the part-

[9]Hires defined as individuals not employed in the same firm in the year prior, excluding switches due to firm ownership changes.

time workers, I exclude workers who were employed in more than one job simultaneously.[10] Finally, as the dynamics of AI skills may differ between public and private sector, I exclude workers in public sector. A full list of the sample restrictions, along with motivations can be found Section C in the Appendix.

2.3 Matching of register data and job vacancy data

The job vacancies serve to provide skill requirements to the actual hires in the register data. The primary objective of this approach is to establish a linkage between individual job vacancies and corresponding hires, to the greatest extent feasible.[11]

The process of matching vacancies to hires is executed through four distinct criteria: the vacancy and the hire must be associated with the same firm, located in the same geographical area (municipality), fall under the same occupation (3 and 4-digit ISCO), and occur within a time window between the posting date and hiring date.

As job advertisements typically do not explicitly mention the intended starting date of the position, the time window between vacancy posting and hiring dates needs to be established. To establish how many month typically pass between vacancy posting and hiring, the following LPM regression is conducted at the firm-location-occupation-month level.[12]

$$\text{Hire}_{flot} = \sum_{k=0}^{12} \gamma_k \cdot \text{Vacancy}_{flo,t-k} + \alpha_t + \epsilon_{flot} \tag{1}$$

Where the dependent variable is an indicator indicating the occurrence of at least one hire.

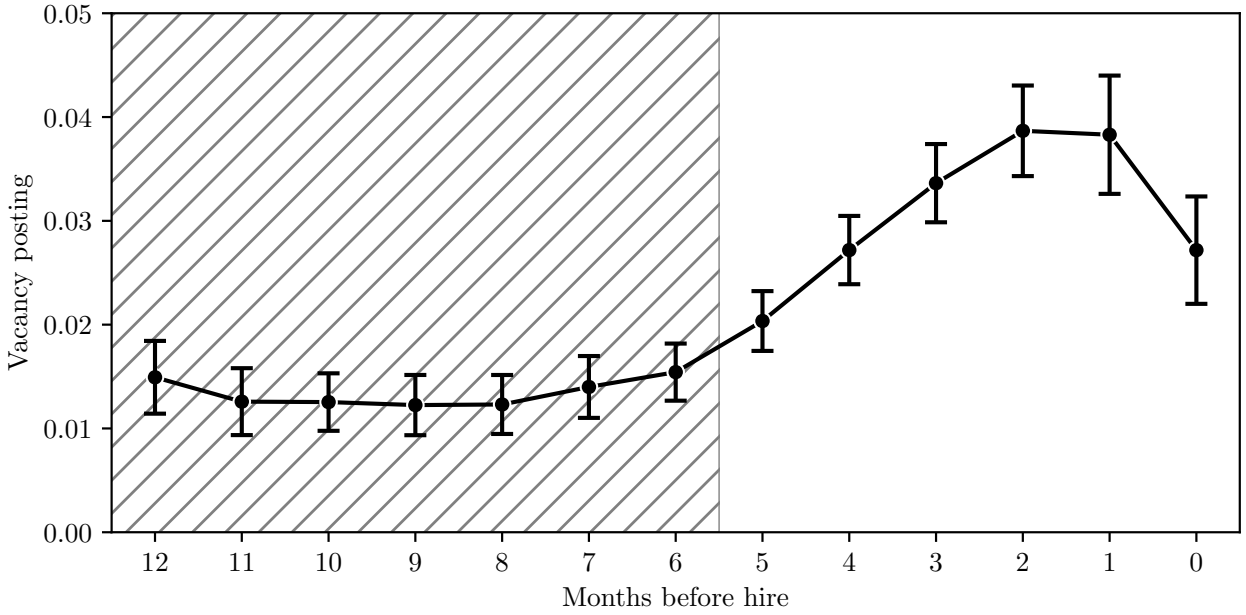
[10] Several of the other sample restrictions has the secondary purpose of reducing the number of part-time workers. Remaining doubts about the impact of these workers are addressed with robustness checks.

[11] A similar matching approach is done by Jensen (2023). As the focus on this paper requires more precise identification of individuals with a rare skill, the matching approach is developed independently.

[12] In the case of later years (2019 onwards), this process is straightforward as monthly data on hiring is available. However, for the remaining sample consisting of yearly data, I approximate the hiring date. The date is approximated using the individuals' share of earnings in the hiring firm in the year of hire, as a share of total earnings in that year, $\hat{T}_{it} = 365\left(1 - \frac{E^{new}_{ift}}{E^{tot}_{it}}\right)$. Full details on this hiring date approximation can be found in Section D in the Appendix. Yearly matching is possible, but reduces precision due to more false vacancy-hire matches, and does not capture cases of vacancies in the end of a year resulting in hires in the following.

The explanatory variables comprise a set of indicators for a matching job vacancy being posted in the 12 months (+ current month) leading up the period, with month fixed effects added to account for overall heterogeneity in vacancy or hiring activity.

FIGURE 2
Hiring period regression



Notes: This figure displays the coefficients of a firm-occupation-month level regression used to determine the time window between posting a vacancy and making a hire, as described in Equation 1. The dependent variable is an indicator variable denoting a hire in a given observation. Explanatory variable consists of indicator variables for connected (firm-occupation) vacancy hiring in each of the 13 months leading up to and including the current month. The regression controls for year-month fixed effects. The non-hatched area represents the period ultimately considered as the recruitment period (current plus 5 months). The error bars represent 95% confidence intervals based on occupation (4-digit ISCO) clustered standard errors.

The graph in Figure 2 displays the result of the regression. In line with a similar exercise on Danish data by Bagger, Fontaine, Galenianos, and Trapeznikova (2022), the results indicate that hiring is most strongly associated with vacancy postings within the five months preceding the hiring month, plus the month of the hire.[13] Consequently, for the purpose of matching hires to vacancies, I consider vacancies posted up to five months prior to the hiring month.

Following this, the four identifiers are used to match individuals to the vacancies corresponding to their positions. This methodology in itself does not ensure that vacancies are tied to

[13]The significant but flat relationship after 5 months is not surprising. A firm that hires a given worker is likely to have engaged in hiring of similar workers in the past.

the correct hire, as a singular firm may advertise multiple matching vacancies in a given time period, thereby generating numerous potential matches for a single hire. Given the focus on a relatively rare skill, a rigorous approach is adopted to maximize the precision of the vacancy-to-hire matching. This involves an iterative one-to-one matching process that *'fills the vacancies'*.

Each iteration consists of identifying all one-to-one matches between vacancies and hires, i.e. hires that can only be tied to one vacancy, which also does not match any other hires. For all one-to-one matches, these vacancies are considered *'filled'*, meaning the pairing is confirmed, and both the individual and the vacancy are excluded from the process for future iterations.[14]

When all one-to-one matches are connected and removed, the next iteration starts. In subsequent rounds, more one-to-one matches are revealed, as hires with previously multiple matches now only have single matches following the removal of *'filled vacancies'*. This approach is ran until there are no more one-to-one matches (with the decreasing pool of unmatched vacancies and hires guaranteeing eventual termination). The iteration step is ran twice. First, using 4-digit ISCO as the occupational identifier until there are no matches. Second, the process is ran again, with occupation relaxed to the broader 3-digit ISCO. This modification acknowledges that there may be slight deviations between the formal occupational code of a vacancy and the final employment contract, but with preference for perfectly matching occupations.

This method leaves a number of hires that cannot be resolved to a single vacancy, as they remain matched to multiple vacancies after the algorithm terminates. These hires are disproportionately concentrated in larger firms, since firms with higher hiring and vacancy volumes are less likely to produce one-to-one matches within a given period. As larger firms are also more likely to use AI (see Table A1 in the Appendix), these workers are important to the

[14] In cases of a single job advertisements consisting of multiple vacancies, a matching hire only *'fills'* one of these vacancies, leaving the remaining vacancies in the advertisement open.

analysis of this paper. I therefore retain these hires by assigning each one a probability of having been recruited to a position requiring AI skills based on the skill content of all their matching vacancies. Specifically, for hire i, let V_i denote the set of vacancies to which they remain matched after the iterative algorithm terminates, and let $\text{AI skill}_v \in \{0, 1\}$ be an indicator variable equal to one if vacancy v requires AI skills and zero otherwise. The AI probability is then:

$$\widehat{\text{AI prob}}_i = \frac{1}{|V_i|} \sum_{v \in V_i} \text{AI skill}_v \tag{2}$$

where $|V_i|$ is the number of vacancies matched to hire i. Each matched vacancy is treated as an equally likely candidate for the position filled by hire i, so $\widehat{\text{AI prob}}_i$ can be interpreted as the fraction of matched vacancies requiring AI skills. This is converted to a binary indicator by imposing a cutoff of $\widehat{\text{AI prob}}_i \geq 0.5$:

$$\text{AI skill}_i = \mathbf{1}\left[\widehat{\text{AI prob}}_i \geq 0.5\right] \tag{3}$$

such that a hire is classified as having AI skills if the majority of their matched vacancies require them. For the one-to-one matches of the iterative approach, $|V_i| = 1$, meaning Equation 3 assigns them an AI skill indicator equal directly to AI skill_v of their single matched vacancy. The cutoff is a tradeoff between precision of AI skill_i and reduced bias towards large firms. The main analysis is found to be robust to restricting the sample to one-to-one matches; see Section F.1 of the Appendix.

2.4 Validation of matching process

The number of matched hires and vacancies are displayed in Table 1. About 23% of hired workers were successfully tied to a job vacancy. While this may seem like a somewhat low number, only 17% of hired workers in Sweden report having been hired via a job vacancy

(SCB, nd). Furthermore, the use of job vacancies is less common in private sector, where methods of direct recruitment without a reported deadline is preferred (SCB, 2023b). The unmatched vacancies represents unfilled vacancies as well as hires substantially deviating from the descriptions of their posted vacancies. 42% of AI vacancies are matched to individuals with a strong enough certainty to be considered as having AI skills.

TABLE 1
Hire-vacancy matching numbers

	AI skills	Full sample
Hires	-	1,605,908
Vacancies	17,793	1,105,758
Matched vacancy-hires	23,675	384,022
Hires with AI skill probability $\geq 50\%$	7,448	-
Perfect matches	3,673	164,255

Notes: This table shows the number of vacancies, hires, as well as matched hire-vacancies. Sample restrictions are imposed.

For validation, I employ another source of within-occupation use of AI, namely the Statistics Sweden survey on ICT usage by individuals. In 2024 and 2025, the survey included questions on use of generative AI tools. The survey was conducted on a stratified sample with around 10,000 respondents between the two years. The survey included questions on whether the respondent uses generative AI tools, for what purpose it is used (work, education, or private use), as well as reason for non-use, which includes two the options of lack of knowledge or awareness of generative AI, as well as insufficient skills in the use of generative AI.

Whereas use of generative AI tools represents a lower level of involvement with AI than what AI skills refer to in the context of this paper, it may still give some indication on the success of the assignment of AI skills to workers via capturing workers' interest and knowledge in the topic. Around 1,100 individuals who responded to the survey in 2024 or 2025 appear and are matched to the vacancy-hire data in the 2014-2023 period (7% of respondents). I run a Linear probability model at the individual level to check whether reported AI use and lack of AI skills can be explained by having been hired to a position requiring AI skills. I

capture the AI skill dummy and controls at the year of the hire.[15]

The results in Table 2 shows no connection between AI skills and use of generative AI tools at work. This is not surprising, given that use of generative AI tools is not considered use of AI in the context of this paper.[16] However, there is a significant negative connection between AI skills and reporting a lack of skills or knowledge of generative AI. Workers with AI skills being 10% less likely to report a lack of generative AI tool skills is reassuring, as one would expect workers involved in more advanced machine learning applications to at least possess the awareness and knowledge of how to use these tools.

TABLE 2
AI hires and reported use of generative AI

	Use of generative AI at work			*Lack of skills & knowledge*		
	(1)	(2)	(3)	(4)	(5)	(6)
AI skill	-0.011	-0.086	-0.096	-0.173***	-0.118***	-0.098**
	(0.105)	(0.117)	(0.124)	(0.011)	(0.040)	(0.048)
Obs.	1,149	1,124	1,124	1,149	1,124	1,124
Adj. R^2	-0.001	-0.007	-0.008	0.001	0.010	0.010
Fixed effects:						
Occupation		✓	✓		✓	✓
Year			✓			✓

Notes: This table displays estimates of a Linear probability model at the individual level with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). In columns 1-3, the outcome variable is a dummy variable for reported use of generative AI tools for work related purposes. In columns 4-6, the outcome variable is a dummy indicating non-use of generative AI tools for the reasons of lack of skills and knowledge required to use such tools. Occupation at 3-digit ISCO (148 possible occupations, relaxed from the otherwise standard use of 429 4-digit to preserve variation with limited number of observations.). Year fixed effects represents the year of the matched hire. Robust standard errors in parantheses. *p<0.1; **p<0.05; ***p<0.01.

3 Workers with AI skills

The matched data consists of hires, of whom some are attributed AI skills. This section presents basic background characteristics, followed by analysis into which observable characteristics are related to sorting into AI occupations, as well as whether there are traits that sets workers with AI skills apart from their occupational peers.

[15] Most recent year of hire in cases of individuals appearing multiple times in the data.

[16] Whereas ICT workers is one of the top users of generative AI tools at work, along with Finance workers and teachers, they only make up 10% of generative AI users, whereas they make up 62% of AI workers.

Table 3 presents basic characteristics of the AI hires.[17] It is revealed that roughly 70% of AI hires are male, and they tend to have higher education than the average hire. Related to the question of returns to AI skills, the earnings of AI hires also stand out immediately. When not accounting for confounders, AI workers typically have higher earnings compared to their non-AI counterparts, with the mean being about 60% higher. Furthermore, AI professionals tend to be younger on average, although they are not predominantly found in the youngest age groups. This can likely be attributed to their higher levels of education, which would require additional years of study before joining the workforce.

TABLE 3
AI worker characteristics

	AI hires	Other hires
Avg. Earnings$_{t+1}$	733.5	439.9
Avg. Age	34.6	33.6
Avg. Experience	10.1	10.3
Share tertiary education	80.4%	48.1%
Share female	30.5%	46.9%

Notes: Notes: This table shows the mean of the characteristics of AI workers (AI probability $\geq$ 50%), and other workers matched to job vacancies. Earnings reported as annual earnings in thousands of SEK.

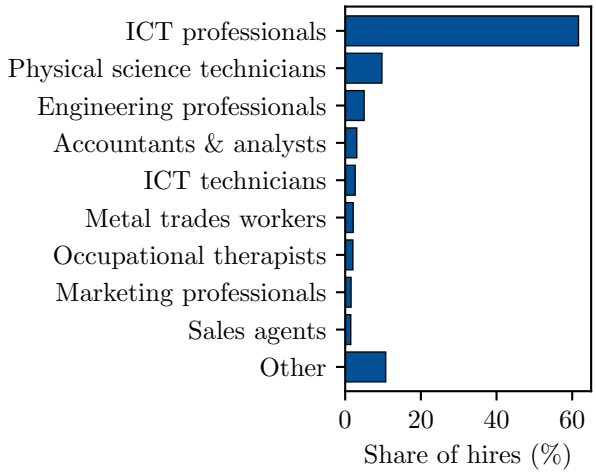


FIGURE 3
AI occupations

Notes: This figure displays the most common occupations of hires with AI skills. Occupations are not controlled for.

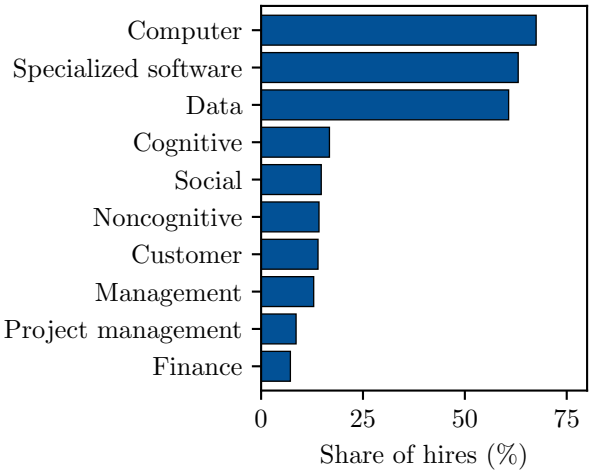


FIGURE 4
AI worker skills

Notes: This figure displays the Deming and Kahn (2018) skills demanded alongside AI skills.

[17] Full descriptives of the different samples used can be found in Table A2 in the Appendix.

This higher education and earnings align with the predominant occupations of AI workers, primarily concentrated in various forms of ICT related occupations, as detailed in Table 3. Most notably, AI workers are highly concentrated within the occupation group of ICT architects[18], although they are found in most occupations where tasks include data analysis or programming.

Given this occupational distribution, the most prevalent skills among AI workers, as depicted in Figure 4, are consistent with expectations. Technical skills such as database and specialized software skills are some of the most common skills of these workers, reflecting the demands of their roles in software development and related domains. In addition to technical expertise, the most crucial soft skills for AI workers include creativity and cognitive skills.

In examining the attributes of firms engaging in the recruitment of individuals possessing AI skills, Table A1 in the Appendix reveals that these firms are substantially larger than the average firm, evident both in terms of their recruitment activities and overall employee count. The dataset identifies 960 distinct firms that have hired at least one worker associated with an AI job posting, hiring on average 1.6 of these workers per year. The industry distribution of AI hires shows that AI hires are most prevalent in the information and communication sector, as well as the finance and insurance sector; see Figure A4 in the Appendix.

3.1 Are AI workers a distinct group?

As seen in Section 3, workers with AI skills overwhelmingly employed in various forms of ICT related occupations. This section investigates whether observable traits can predict which workers sort into AI occupations, and whether there are traits predicting AI workers as compared to their occupational peers. This can partly imply whether AI workers should be considered a subset of AI pioneers within occupations, or whether all ICT workers are

[18] While a large share of AI hires are ICT architects, there is still within-occupation variation; only 9% of all hired ICT architects are AI hires.

potential AI workers, if employed in a position where their AI skills are utilized.

I investigate five potential observable traits that may separate AI workers from their peers. Namely, education, ability, AI experience, age, and migration.

Taking inspiration from the descriptive approach of Bell, Chetty, Jaravel, Petkova, and Van Reenen (2018), I first compare the distribution of workers over the variables of interest between workers hired into positions requiring AI skills, versus the ones who do not. In order to disentangle workers from their occupation, I weight the shares of non-AI hires to match the occupational composition of AI hires. As AI hiring increases over time, this approach risks capturing correlations via time trends if year of hiring is not accounted for. I therefore extend the weighting to be Occupation#year based.

Second, to provide evidence on the key traits defining AI workers and AI occupations, Section 3.1.4 provides an econometric decomposition of AI skills, occupational sorting, and within-occupation differences.

3.1.1 Education

It is possible that while AI workers are hired as ICT occupations, their educational background has provided them a distinct skill set that separates them from their occupational peers. Figure 5 reveals that AI workers are significantly more likely to have a university degree than the average workers. However, when comparing the workers to the reweighted comparison group, it is clear that this difference is driven by the occupational composition of these workers. Comparing AI hires to their reweighted counterparts, AI workers seem to have slightly higher levels of university education, especially in terms of workers with doctorate degrees.

The bottom panel of the Figure displays tertiar edcuation subjects, split by STEM and Non-STEM subjects.[19] STEM subjects are further divided by the main 1-digit ISCED

[19]STEM subjects are identified as being related to AI technologies by Bonfiglioli et al. (2024); Humlum and Bjørnsson Meyer (2022)

2011 subject codes 07 'Sciences' (Mathematics, statistics, physics, etc.), 08 'ICT' (Computer science), and 09 'Engineering'.

The subjects studied by AI hires tell a similar story. AI related occupations are much more likely to have various STEM degrees, but AI hires are only slightly more likely to have a degree in each of the STEM subfields when compared to reweighted non-AI workers.

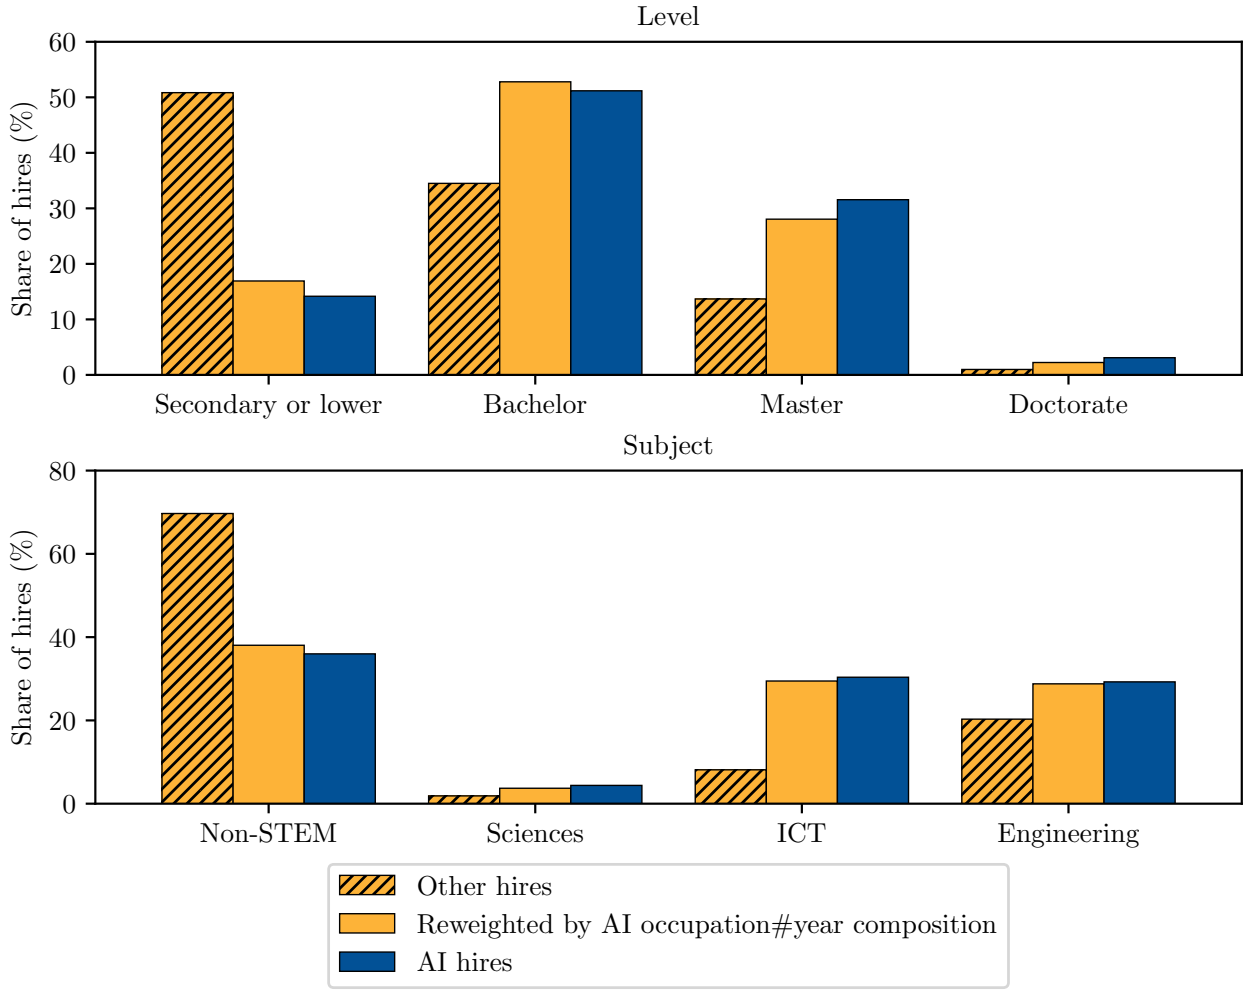


FIGURE 5

AI hire education level

Notes: These figures illustrate the distribution of hires among education levels (top), as well as subjects (bottom). Education subjects are based on the subfields of STEM, ISCED 2011 subjects 07 (Science), 08 (ICT, includes ICT related civil engineer subjects 523a, 523c, 523e, and 523x from the Swedish educational classification SUN 2000, which otherwise fall under the ISCED 2011 engineering subject 09), and 09 (Engineering). The blue bar (right) represents workers matched to vacancies requiring AI skills. The yellow hatched bar (left) represents other hires. The yellow solid bar (middle) represents other hires reweighted by occupation#year AI worker prevalence. The shares sum up to 96 (unweighted), 93 (weighted), and 94 (AI workers). Remaining workers include education lower than secondary, as well as missing educational information.

3.1.2 Ability

Despite AI workers not standing out clearly in terms of education, there could also be other, less observable differences that enable these workers to take a leading role in driving the uptake of new technologies.

For cognitive ability, there are two sources of data available. First, I use math grades from the ninth year of mandatory primary school (Student ages 15-16).[20]

Second, I use the results of military conscription cognitive tests. The data covers individuals conscripted in the period 1997-2024. As part of the enlistment, all conscripted undergo a test in cognitive ability.[21] This rated the conscripted on a scale 1-9 in reasoning, verbal comprehension, spatial ability, and technical understanding. Conscription was mandatory for men until 2010, when it was abolished. Mandatory conscription was reinstated in 2017, this time gender-neutral. There is still test data from conscripts in the period in-between mandatory conscription, albeit with lower numbers (About 25% the levels of mandatory years.). Data is available for 18% of the hire sample for men, and 2% of women. Following the approach of Nybom, Plug, van der Klaauw, and Ziegler (2023), I increase coverage via imputation from relatives. I impute test score for missing values by the mean test score of siblings via the father. This increases the coverage to 27% of men, and 17% of women.

I employ the math grades and conscription test scores complementary. The math scores are available for most of the population, but may be biased by socioeconomic status. The conscription data is less affected by confounders, but with the sample is heavily biased towards men.

[20]The use of math grades as a proxy for cognitive ability is used by Bell et al. (2018). The majority of grades (Period 1998-2012) uses a system of three passing grades. I assign these grades numerical values of 1-3. For periods prior and after the three grade period, there are five passing grades. As these are less common, I transform these grades to the three grade system, with 1-2 corresponding to 1, 3-4 corresponding to 2, and 5 corresponding to 3. The ninth grade is the only year of primary school grades available in the data.

[21]This data has been used by e.g. Grönqvist and Lindqvist (2016); Håkanson, Lindqvist, and Vlachos (2020); Rönnlund, Carlstedt, Blomstedt, Nilsson, and Weinehall (2013)

Both measures show that AI hires have somewhat higher levels of cognitive ability on average as compared to their occupational peers; see Figure 6.

The noncognitive ability may be important not so much in explaining who acquires AI skills, but rather reflect the ability of an individual to market their skills, thereby having explanatory power on who gets hired into positions requiring certain skills.

Figure 7 shows that similarly to cognitive ability, AI workers and reweighted peers have higher noncognitive ability. Among the higher than average grades 6 and 7, the AI workers stand out from their peers with higher scores.

3.1.3 AI experience and age

Another possible way that AI hires may acquire a unique skill set is by working with AI in the past. To check this, I construct a measure of potential AI experience. This is a measure of years that an individual have worked in a firm with confirmed AI use, in an occupation related to AI. Firms using AI is defined as any time period after and including a firm has posted at least two AI related vacancies. AI occupations are defined as the occupations where 90% of AI vacancies are posted, sorted by occupation share of AI vacancies.

The AI experience measure is shown in figure 8. As with the prior variables, there is a hint that AI hires are slightly more likely to have worked in an AI related role in the past. However, it is once again not a defining feature, with mostly overlapping backgrounds in comparison to peers.

Given that AI skills reflect cutting edge technologies, one may assume that workers who studied a more recent curriculum would be more likely to acquire AI skills. The distribution of ages show that AI workers are somewhat younger than their peers, but more concentrated around middle ages than other workers. This reflects the fact that AI workers and their peers have much higher tertiary share than the average worker, making them older by the time they enter the labor force for the first time.

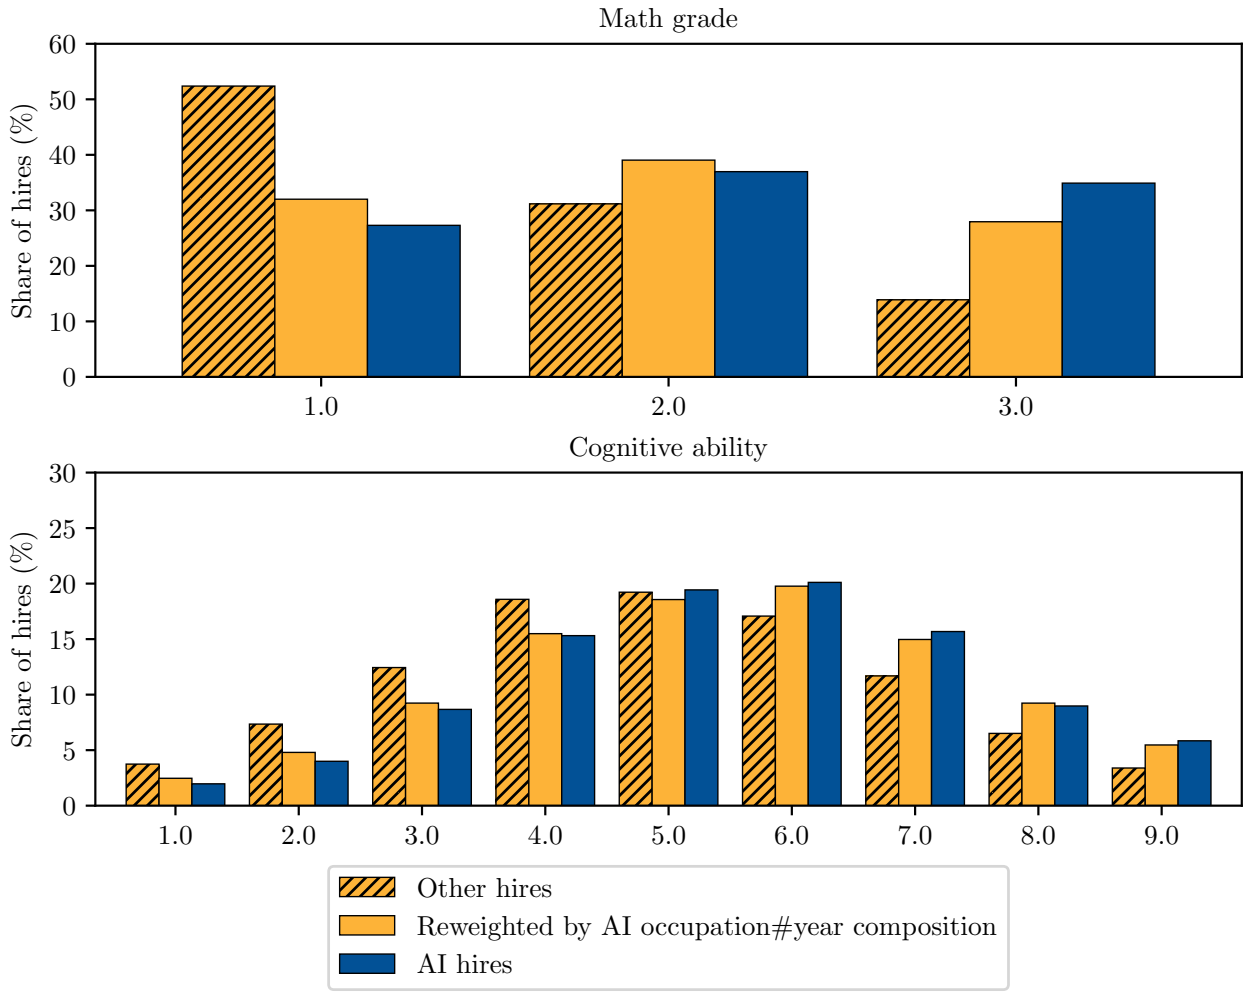


FIGURE 6

AI hire cognitive ability

Notes: This figure illustrates the distribution of hires among ability levels, proxied by math grades from 9th grade of primary school, ages 15-16 (top panel), as well as by the national conscription test data (bottom panel). Missing data imputed by mean sibling ability where available. The blue bar (right) represents workers matched to vacancies requiring AI skills. The yellow hatched bar (left) represents other hires. The yellow solid bar (middle) represents other hires reweighted by occupation#year AI worker prevalence.

A final possibility is that AI workers are not primarily supplied by the Swedish economy, but are rather recruited from abroad. This is an important complement to the prior measures, as the data on education, grades, conscription, and AI experience would not cover recent immigrants. To check this, I consider hires who immigrated to Sweden within the three years prior to the hire (including the year of hiring.). Figure 10 shows that workers in ICT occupations are twice as likely to be recent immigrants, these workers are not more likely to bring specialized AI expertise.

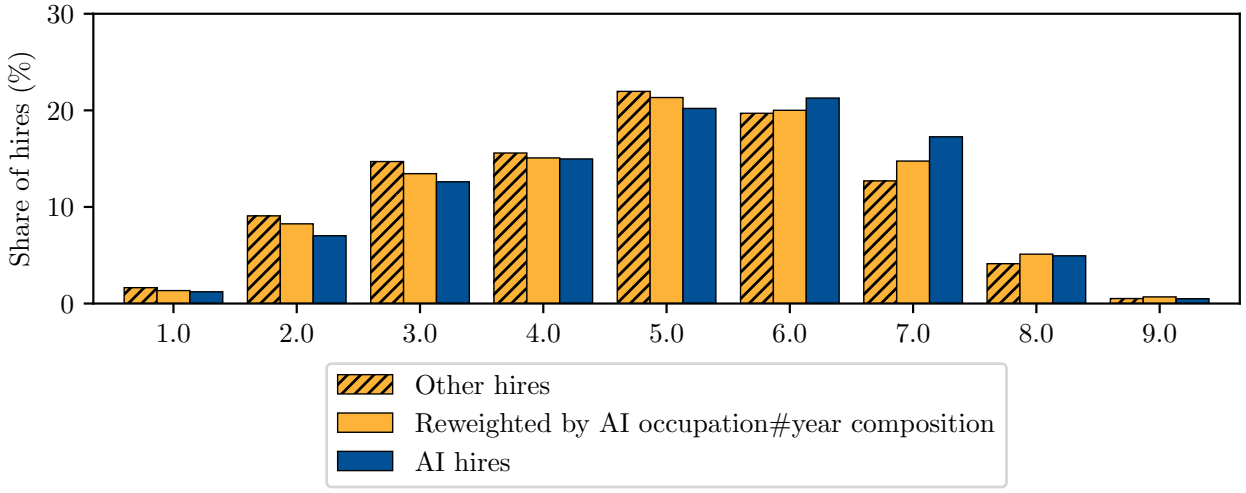


FIGURE 7
AI hire noncognitive ability

Notes: This figure illustrates the distribution of hires among ability levels, proxied by math grades from 9th grade of primary school, ages 15-16 (top panel), as well as by the national conscription test data (bottom panel). Missing data imputed by mean sibling ability where available. The blue bar (right) represents workers matched to vacancies requiring AI skills. The yellow hatched bar (left) represents other hires. The yellow solid bar (middle) represents other hires reweighted by occupation#year AI worker prevalence.

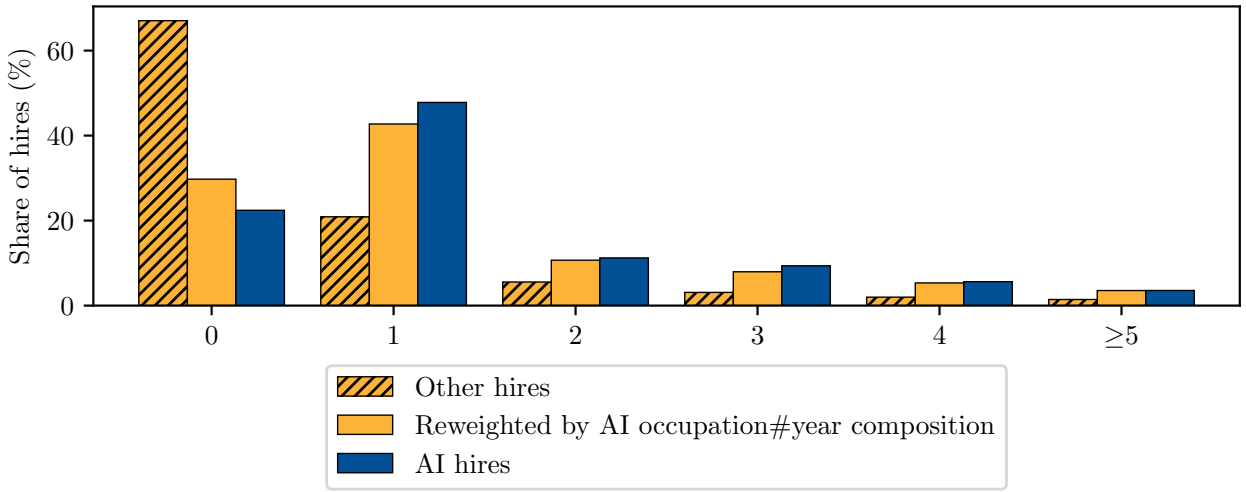


FIGURE 8
AI hire past AI experience

Notes: This figure illustrates the distribution of hires among years of potential AI experience. Potential AI experience is defined as years employed in a firm that posted at least two AI vacancies before the year of hire, in the occupations where 90% of AI skill demand is concentrated (in order of highest share of AI vacancies.). The blue bar (right) represents workers matched to vacancies requiring AI skills. The yellow hatched bar (left) represents other hires. The yellow solid bar (middle) represents other hires reweighted by occupation#year AI worker prevalence.

3.1.4 Overall evidence

The descriptives point towards most differences in AI workers and the average worker coming from the occupational sorting. In relative terms, workers with AI skills only display slightly

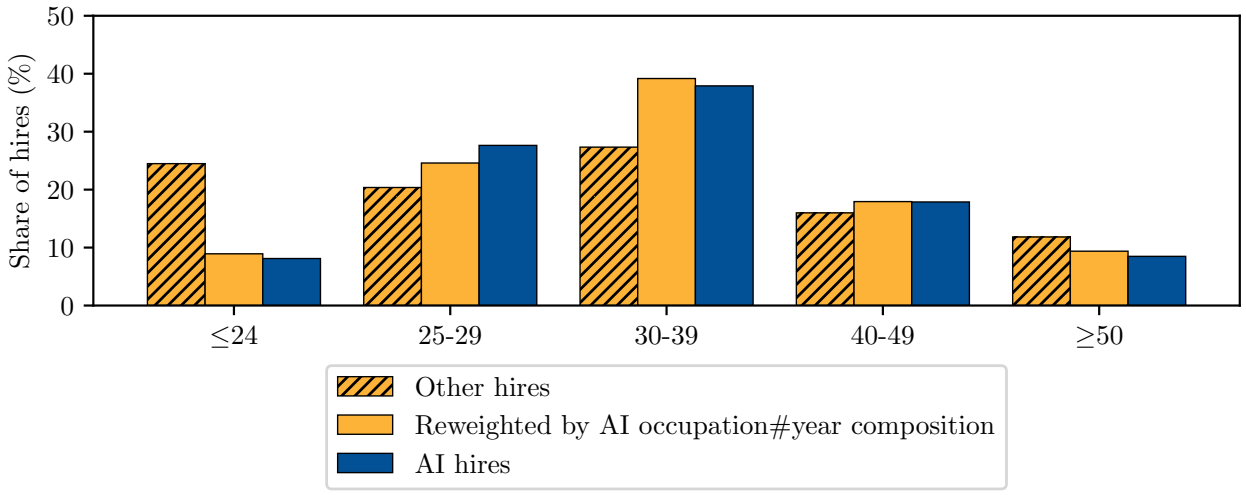


FIGURE 9
AI hire age

Notes: This figure illustrates the distribution of hires among age groups. The first age brackets split by 5 year intervals to separate ages typical for tertiary graduation. The blue bar (right) represents workers matched to vacancies requiring AI skills. The yellow hatched bar (left) represents other hires. The yellow solid bar (middle) represents other hires reweighted by occupation#year AI worker prevalence.

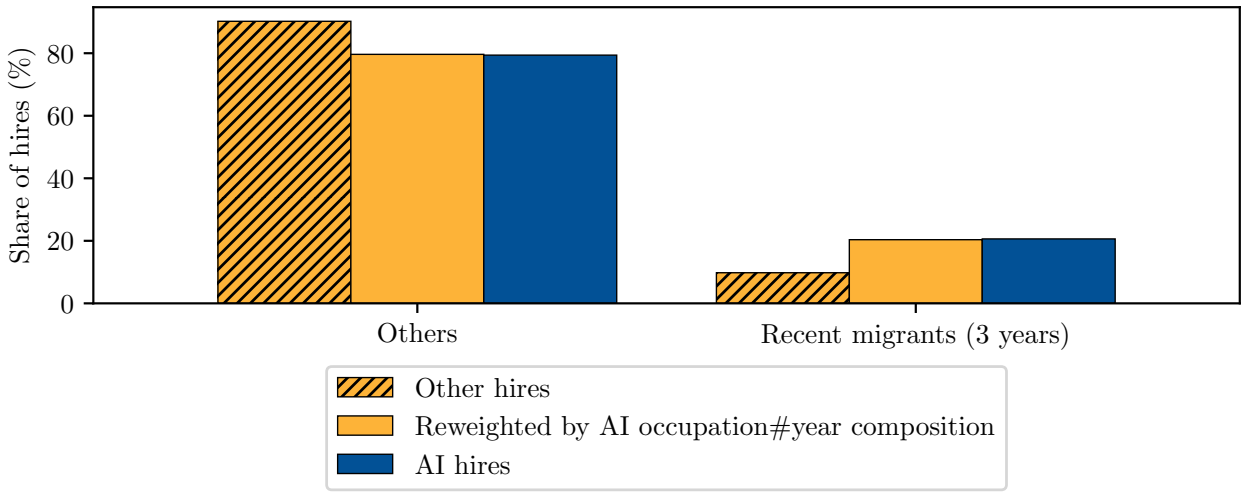


FIGURE 10
AI hire recent migration

Notes: This figure illustrates the distribution of hires among workers who migrated to Sweden in the last three years (including year of hiring), and others. The blue bar (right) represents workers matched to vacancies requiring AI skills. The yellow hatched bar (left) represents other hires. The yellow solid bar (middle) represents other hires reweighted by occupation#year AI worker prevalence.

higher levels of education, cognitive ability, and AI experience, compared to their peers in the same occupations. I employ a within-between model of Bell and Jones (2015); Mundlak (1978) to disentangle the observable traits that are generally associated with workers in occupations using AI (Occupation sorting), and the traits that separate workers from their

occupational peers (Within-occupation).[22]

Table 4 reports estimates for education. Occupations with higher levels of AI use tend to have higher levels of education. Doctorate degrees stand out as having a greater connection to AI skills than other education levels. Workers with AI skills are also likely to have a higher level of education. Unsurprisingly given the occupations AI workers consist of, workers with AI skills are more likely to have a degree in ICT or Sciences. In comparison to occupational peers, workers with AI skills only display small associations with higher education levels, and are slightly more likely to have an education within sciences.

TABLE 4
AI occupation & AI worker education

	Degree (1)			Subject (2)		
	Bachelor's	*Master's*	*Doctorate*	*Sciences*	*ICT*	*Engineering*
Occupation sorting	0.065***	0.054**	0.144*	0.064**	0.222***	0.007
	(0.024)	(0.027)	(0.085)	(0.026)	(0.015)	(0.009)
Within-occupation	0.002**	0.008***	0.016***	0.007**	0.003	0.001
	(0.001)	(0.002)	(0.006)	(0.004)	(0.002)	(0.002)
Occ. sorting share	97.7%	87.4%	89.9%	90.0%	98.7%	84.1%
Obs.		384,028			384,028	
Adj. R^2		0.020			0.050	

Notes: This table displays estimates of two separate individual-hire level Linear Probability Model with an AI skill dummy (1 = AI probability $\geq$ 50%) as the outcome variable. In columns labeled (1), the explanatory variables are levels of education (each degree includes workers with that level or higher.). In columns labeled (2), the explanatory variables are STEM education subjects. Subjects divided by the main 1-digit ISCED 2011 subject codes 07 'Sciences' (Mathematics, statistics, physics, etc.), 08 'ICT' (Computer science), and 09 'Engineering'. The regression is a a within-between occupation setup $AIskill_{iot} = \beta_0 + \beta_1(x_{iot} - \bar{x}_{ot}) + \beta_2\bar{x}_{ot} + \epsilon_{iot}$. Occupation sorting displays occupational traits associated with AI skills ($\beta_2\bar{x}_{ot}$). Within-occupation represents the traits associated with AI skills beyond occupational traits ($\beta_1(x_{iot} - \bar{x}_{ot})$). The Occ. sorting share reports the share of the relation between the trait and AI skills attributed to the sorting into AI related occupations, $\frac{\hat{\beta_2}}{\hat{\beta_1}+\hat{\beta_2}} * 100$, omitted in cases of opposite signs. The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation level. *p<0.1; **p<0.05; ***p<0.01.

Table 5 displays results for worker traits. Cognitive ability, past AI experience, as well as being a recent immigrant can predict AI skills in a worker. Cognitive ability does not predict AI skills beyond sorting into AI occupations. Workers with AI experience are, however, more likely to use AI than their peers, but it only explains around 8% of the relation between AI experience and AI skills.

[22] $AIskill_{iot} = \beta_0 + \underbrace{\beta_1(x_{iot} - \bar{x}_{ot})}_{\text{Within-occupation}} + \underbrace{\beta_2\bar{x}_{ot}}_{\text{Occupation sorting}} + \epsilon_{iot}$, where i is denotes individuals, o occupation (4-digit ISCO), t year. The regression provides identical estimates to an occupation#year fixed effect model $AIskill_{iot} = \beta_0 + \beta_1 x_{iot} + \gamma_{ot} + \epsilon_{iot}$, but with reported coefficients for the occupation sorting component.

TABLE 5
AI occupation & AI worker traits

	Ability (1)		Background (2)		(3)
	Cognitive Ability	*Noncognitive Ability*	*AI Experience*	*Age*	*Recent migration*
Occupation sorting	0.051***	-0.001	0.032***	-0.023*	0.203***
	(0.020)	(0.007)	(0.008)	(0.014)	(0.072)
Within-occupation	0.001	0.001	0.003**	-0.001**	-0.001
	(0.000)	(0.001)	(0.002)	(0.001)	(0.005)
Occ. sorting share	98.9%	–	91.2%	94.8%	–
Obs.	68,142		291,221		384,028
Adj. R^2	0.012		0.023		0.022

Notes: This table displays estimates of three individual-hire level Linear Probability Model with an AI skill dummy (1 = AI probability $\geq 50\%$) as the outcome variable. The explanatory variables are predictors of AI skills, in a within-between occupation setup $AIskill_{iot} = \beta_0 + \beta_1(x_{iot} - \bar{x}_{ot}) + \beta_2\bar{x}_{ot} + \epsilon_{iot}$. Occupation sorting displays occupational traits associated with AI skills ($\beta_2\bar{x}_{ot}$). Within-occupation represents the traits associated with AI skills beyond occupational traits ($\beta_1(x_{iot} - \bar{x}_{ot})$). All explanatory variables bar recent migration are standardized ($\tilde{x} = \frac{x-\mu_x}{\sigma_x}$). Coefficients represent the change in the probability of having an AI skill associated with a one standard deviation increase in x ($\Delta x = \sigma_x$), corresponding to the following values: Math Grade = 0.76, Cognitive Ability = 1.93, Noncognitive Ability = 1.72, AI Experience = 1.05, Age = 11.16. The Occ. sorting share reports the share of the relation between the trait and AI skills attributed to the sorting into AI related occupations, $\frac{\hat{\beta}_2}{\hat{\beta}_1+\hat{\beta}_2} * 100$, omitted in cases of opposite signs. The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. The number of observations in regression (1) is lower due to data availability of conscription cognitive tests. Regression (2) excludes ages ≤ 24 due to the high prevalence of university education leading to these workers being older when first entering the workforce. AI experience coefficients similar without the age restriction. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation level. *p<0.1; **p<0.05; ***p<0.01.

Taken together, Cognitive ability, having a degree in ICT subjects, and being a recent immigrant explain sorting into AI occupations, whereas only high-level education, and in sciences subjects, together with past AI experience can explain some within-occupation AI prevalence. Throughout, the within-occupation coefficients are small to the point of not being economically meaningful. The largest coefficient rules out a doctorate degree being associated with a probability of having AI skills larger than 3% beyond their occupational peers.

According to these observable traits, workers in AI related occupations should be seen as potential AI workers, and there is no significant evidence pointing to a small pioneering group within these occupations responsible for pushing the technology. Next section turns to potential unobservable traits by investigating wage dynamics of workers with AI skills.

4 Returns to AI skills

The previous section found that while observable characteristics such as cognitive ability, AI experience, as well as educational background can predict which workers enter AI occupations, there is only weak evidence towards observable characteristics explaining who becomes an AI worker beyond the occupational sorting. There is, however, still the possibility that the observable characteristics do not fully capture all relevant aspects of AI occupations. Earnings can be seen as a proxy for the value of the worker, in which unobservable characteristics are bundled. A significant earnings premium of AI skills when controlling for observable traits would indicate that there are unmeasured characteristics valued by firms, beyond what is typical for their peers.

The basis for the empirical analysis is the data of new hires matched to skill content of vacancies.[23] The baseline regression uses the model in Equation 4.

$$ln(\text{earnings}_{if,t+1}) = \beta_0 + \beta_1 \text{AI skill}_{ift} + \boldsymbol{\beta_x X_{it}} + \boldsymbol{\beta_z Z_{ft}} + \epsilon_{ift} \tag{4}$$

Where t is the year of hire of individual i in firm f. $earnings_{if,t+1}$ is the annual earnings of the individual in the year following the hire. $t+1$ is used to get a representative full year of earnings in the new job. $\boldsymbol{X_{it}}$ is a vector of individual control variables, and $\boldsymbol{Z_{ft}}$ is a vector of firm control variables, both taken from the year of the hiring. $AIskill_{ift}$ is a dummy indicating whether the hired individual is matched to a vacancy requiring AI skills.

Unless otherwise stated, all regressions include control variables for gender, age, work experience, and experience squared. In addition, fixed effects are integrated for year, location[24], age of youngest child, and the interaction of female gender with the age of the youngest child.

[23] The choice of only including individual-years at the time of the hire is motivated by the fact that job switching in itself tend to lead to high earnings increases, which makes the relevant control group other job switchers.

[24] FA-15 municipality groups, similar to commuting zones.

As education level, having a degree in sciences subjects, as well as past AI experience are all found to be significantly linked to AI skills within-occupation, how these relate to a potential AI earnings premium is important both in terms of separating an earnings premium from the observables, but also by seeing how AI skills are valued in different settings. As such, regressions are split by excluding these 'AI observables', controlling for them, and interacting them with AI skills.

TABLE 6
AI earnings premium

	Base premium			*Controls for observables*			*Observables interactions*		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
AI skill	0.419***	0.054***	0.009	0.264***	0.043***	0.003	0.429***	-0.039	-0.087**
	(0.051)	(0.011)	(0.011)	(0.043)	(0.011)	(0.011)	(0.052)	(0.044)	(0.043)
AI experience				0.099***	0.040***	0.036***	0.101***	0.040***	0.036***
				(0.018)	(0.005)	(0.005)	(0.019)	(0.005)	(0.005)
ai_x_ai_experience							-0.057***	0.001	0.008
							(0.019)	(0.009)	(0.008)
AI * Science degree							-0.028	-0.087	-0.066
							(0.097)	(0.081)	(0.073)
Bachelor * AI							-0.068	0.112***	0.110**
							(0.064)	(0.043)	(0.046)
Master * AI							-0.191***	0.087**	0.084**
							(0.054)	(0.044)	(0.042)
Doctorate * AI							-0.197***	0.165***	0.189***
							(0.049)	(0.045)	(0.051)
Obs.	278,564	278,564	278,564	278,564	278,564	278,564	278,564	278,564	278,564
Adj. R^2	0.190	0.298	0.339	0.222	0.301	0.340	0.222	0.301	0.340
Control variables:									
Basic individual variables	✓	✓	✓	✓	✓	✓	✓	✓	✓
Observable AI traits				✓	✓	✓	✓	✓	✓
Fixed effects:									
Occupation		✓	✓		✓	✓		✓	✓
Workplace industry			✓			✓			✓
Firm			✓			✓			✓

Notes: This table displays estimates of a regression with the natural logarithm of earnings as the outcome variable, with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). Columns 4-6 controls for average lifetime earnings before year t, as well as years of experience interacted with average lifetime earnings. The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. All regressions use the most restrictive sample. The number of observations in columns 4-6 is lower due to individuals with no recorded wages in t_{-1} being excluded. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation and firm level. *p<0.1; **p<0.05; ***p<0.01.

Table 6 presents the findings of the baseline AI skill earnings premium regression. The table reveals a 5% premium compared to peers in the same occupation[25] in the Appendix displays additional variations of the controls, revealing a a premium of 22% when not controlling for

[25]Earnings premium calculated as $(e^{\hat{\beta}_1} - 1) * 100$ in regressions using $ln(earnings_{if,t+1})$ as the outcome. See Table **??** for additional specifications

occupation. Similar to the observable trait regressions in Section 3.1.4, the occupation of the worker can explain 90% of the earnings differences between AI and non AI workers. However, the coefficient goes nonsignificant when adding firm fixed effects in column 3, implying that firms hiring AI workers tend to pay higher salaries.

The earnings premium is seemingly only explained by confounding observable AI traits to a small degree, with results being largely unchanged when adding these controls.

However, interacting the AI skill dummy with the observable AI traits reveals some interesting and important dynamics. While experience in AI related roles is valued by employers, more experienced AI hires do not acquire higher earnings. A degree in science subjects is also not related to higher earnings. In terms of education, the negative coefficient of AI skills is balanced out by the coefficients of bachelor's and master's degrees. However, AI workers with a doctorate degree earn 20% more than non-AI doctorates, and 9-12% more than tertiary educated AI workers. This result importantly points towards a split among AI workers, where AI users are mostly explained by their occupations, there is a small and highly valued subsample of AI developers who stand out beyond their peers.

The previous result suggests that higher wages experienced by AI hires are explained by higher wages of AI using firms. To scrutinize this, I investigate earnings at the firm-level directly. These regressions use mean firm wages as the outcome, explained by whether the firm posted any AI vacancies in the 2014-2023 period.

The result in Table 7 shows that firms using AI tend to pay their workers 11% more when controlling for size differences, industry, and occupational composition of firms. While part of this may stem from spillovers from use of AI, a more likely explanation is that high paying firms tend to be the ones to use AI.

TABLE 7
Wages of AI using firms

	(1)	(2)	(3)
AI hiring(1 = min. 1 hire)	0.519***	0.148***	0.111***
	(0.054)	(0.020)	(0.017)
Obs.	17,669	17,669	17,388
Adj. R^2	0.037	0.510	0.560
Fixed effects:			
Firm size		✓	✓
Industry		✓	✓
Occupation			✓

Notes: This table displays estimates of a regression at the firm level with mean firm wages in the period 2014-2023 as the outcome variable, with the explanatory variable being a dummy representing at least one AI hire in the period 2014-2023(1 = AI probability $\geq$ 50%). All control variables taken from year 2014. Sample is limited to firms present in the main regression in Table **??**. All regressions use the most restrictive sample. Industry as fourth-level NACE. Occupation denotes occupational composition controls represented by variables for firm employment share of each 4-digit ISCO. Standard errors in parantheses are clustered at industry level. *p<0.1; **p<0.05; ***p<0.01.

4.1 Entering an AI job

There are multiple arguments both for and against why entering a position requiring AI skills may affect the earnings growth of an individual. Increased earnings may be expected either if an individual developed their skills recently, as well as in cases where the AI skills of an individual were not fully utilized in their previous position. Further, the basic earnings premium found in Section 4 found that AI skills is associated with employment in high-paying firms. These arguments all make assumptions about the background of workers with AI skills. It is not unlikely that these individuals worked with AI in past roles, and may also have been employed in high-paying firms previously.

To estimate the earnings premium associated with transitioning into a position requiring AI skills, Callaway and Sant'Anna (2021) event study is employed. The event is set as the year of the job switch, with the panel limited to the stint before and after the switch. Individuals hence include one panel for every vacancy-matched job change. As switching jobs is associated with high increases in salary regardless of the position one transitions into, I compare the estimates of switching into jobs requiring AI skills, to a separate event study

including all job switches. As this setup would violate the parallel trends assumption, I restrict the control group to no-yet-treated individuals.[26]

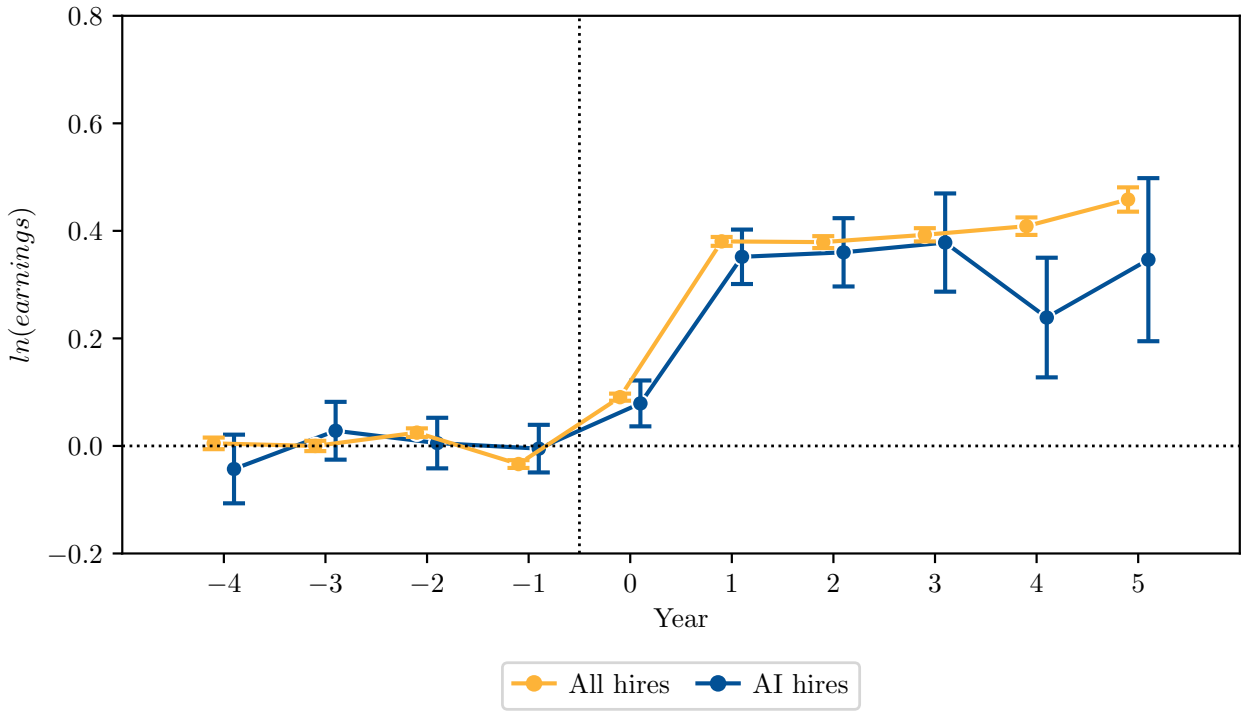


FIGURE 11
AI earnings event study

Notes: This figure displays estimates of a the event study regression described in Equation **??**, with the natural logarithm of earnings as the outcome variable. The dark blue line represents estimated earnings around getting hired into any new job. The light yellow bar represents the estimates for being hired into the subset of hires for positions requiring AI skills (1 = AI probability $\geq$ 50%). The sample is restricted to hires matched to job vacancies with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. The error bars displays a 95% confidence interval based on standard errors clustered at individual level.

The result of the event study is displayed in Figure 11. Whereas job switching is associated with substantially increased earnings, entering a job requiring AI skills does not stand out from the average job switch.

This result implies that the unobserved characteristics leading to higher earnings of workers with AI skills may be more persistent traits, such as personal characteristics including an interest in new technologies, as well as unobserved skills acquired in the past.

These event studies must recognize multiple key limitations. First, Entering a job is a

[26]The hypothetical control group would either be individuals who never switched jobs, or individuals who switched to non-AI jobs, neither of which would be a desirable control group.

choice made by the individual, and not an exogenous event. Second, whereas the model accounts for the higher level of earnings experienced by workers with AI skills, there may be multiple characteristics affecting the magnitude of growth in earnings, unrelated to the position requiring AI skills or not. Most crucially, this would be correlated with experience, as more experienced workers would experience diminishing growth in earnings. Lastly, as AI hiring increases over time, the estimates from later post-event years is based on fewer individuals, likely explaining the volatility in years four and five.

4.2 Robustness checks

Robustness checks can be found in Section F in the Appendix. The chosen matching method of using a 50% cutoff for AI skill probability is scrutinized in Section F.1, wherein I find that the results hold when using the perfectly matched sample, as well as any other cutoff value. Subsequently, in Section F.2, I evaluate the potential impact of not controlling for part-time employment, finding similar coefficients. In section E.1, I attempt to disentangle the earnings premium associated with AI skills from the closely intertwined data and software skills. While the effect of AI skills mostly persists when including data and software skills, the effect is not significantly different from these skills. In section F.3, I employ the methodology of the closely related paper by Alekseeva et al. (2021) as closely as possible. This check verifies that the differing methodology is not the reason for the lower premium found in my study. Lastly, in Section F.4, I make sure that the choice of always using the sample from the most restrictive specification of the regressions leads to similar results as using the unrestricted sample.

5 Conclusion

Despite recent policy interest and reported scarcities, not much is known about what constitute AI workers, and what traits are needed in the training of them. I investigate whether AI workers should be seen as a group of occupations related to AI, or a subset of technology

pioneers with the key skills needed to develop and use AI in firms. I employ a method of matching job vacancies to register data on hired individuals to investigate observable as well as unobservable characteristics of workers with AI skills.

I find that the average AI worker are workers in occupations where a tertiary degree in Sciences and ICT subjects, high cognitive ability, past AI experience, and being a recent migrant are common traits.

Beyond the observable traits, I find that workers with AI skills have 5% higher earnings than their peers, a figure notably lower than previously documented in a study of wage premiums in the US labor market (Alekseeva et al., 2021). This is unsurprising, given that the Swedish labor market is characterized by comparatively rigid wages as a result of collective bargaining. This base earnings premium stems entirely from the high wages of the firms posting AI vacancies. As such, workers with AI skills have a higher chance of being matched to a high paying firm, but does not get higher compensation than their peers within those firms.

These two findings indicate that the average AI worker are mostly characterized by the occupation they are employed in, with unobservable characteristics only explaining some additional likelihood of sorting into high paying firms.

However, interactions reveal that within-occupation, AI workers are more likely to have a doctorate degree. AI workers with doctorate degrees are also highly valued by employers, acquiring significantly higher earnings than even their firm and occupation peers.

Taken together, the results suggest that the supply of the average AI worker is closely linked to the supply of STEM graduates. To supply the economy with an elite of AI developers, individuals with research experience are needed. The second group is highly demanded by the private sector.

This has policy implications that vary depending of the goal of the policymaker; To spread basic AI use, policies need to focus on increasing the number of students choosing a STEM

education, or the immigration of these. To supply the workers needed to become a leader in private sector AI development, a supply of doctorates in machine learning is crucial. Private sector co-financing of industrial machine learning PhD students might alleviate the costs of the resulting steady brain drain from academia to private sector.

This wider question of AI adoption would benefit from additional research outside the scope of my paper. The paper abstracts from possible adoption via the outsourcing of AI services. This may be a crucial channel for alternative adoption, and should be investigated further. Additionally, the differences in the traits and training of the average AI worker and the pioneering elite would benefit from additional research into the demand and separate roles of the two groups.

References

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In D. Card & O. Ashenfelter (Eds.), (Vol. 4, p. 1043-1171). Elsevier. doi: https://doi.org/10.1016/S0169-7218(11)02410-5

Acemoglu, D., Autor, D., Hazell, J., & Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, *40*(S1), S293-S340. doi: 10.1086/718327

Acemoglu, D., & Restrepo, P. (2018, January). *Artificial intelligence, automation and work* (Working Paper No. 24196). National Bureau of Economic Research. doi: 10.3386/w24196

Acemoglu, D., & Restrepo, P. (2019, May). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, *33*(2), 3-30. doi: 10.1257/jep.33.2.3

Alekseeva, L., Azar, J., Giné, M., Samila, S., & Taska, B. (2021). The demand for ai skills in the labor market. *Labour Economics*, *71*, 102002. doi: https://doi.org/10.1016/j.labeco.2021.102002

Alekseeva, L., Gine, M., Samila, S., & Taska, B. (2020). *Ai adoption and firm performance: Management versus it.* Retrieved from https://ssrn.com/abstract=3677237

Babina, T., Fedyk, A., He, A. X., & Hodson, J. (2023). *Firm investments in artificial intelligence technologies and changes in workforce composition* (Vol. 31325). National Bureau of Economic Research.

Bagger, J., Fontaine, F., Galenianos, M., & Trapeznikova, I. (2022). Vacancies, employment outcomes and firm growth: Evidence from denmark. *Labour Economics*, *75*, 102103. doi: https://doi.org/10.1016/j.labeco.2021.102103

Baruffaldi, S., van Beuzekom, B., Dernis, H., Harhoff, D., Rao, N., Rosenfeld, D., & Squicciarini, M. (2020). *Identifying and measuring developments in artificial intelligence* (Working Paper No. 2020/05). OECD Publishing. doi: https://doi.org/https://doi.org/10.1787/5f65ff7e-en

Bell, A., Chetty, R., Jaravel, X., Petkova, N., & Van Reenen, J. (2018, 11). Who becomes an inventor in america? the importance of exposure to innovation*. *The Quarterly Journal of Economics*, *134*(2), 647-713. Retrieved from https://doi.org/10.1093/qje/qjy028 doi: 10.1093/qje/qjy028

Bell, A., & Jones, K. (2015). Explaining fixed effects: Random effects modeling of time-series cross-sectional and panel data. *Political Science Research and Methods*, *3*(1), 133–153. doi: 10.1017/psrm.2014.7

Bloom, N., Hassan, T. A., Kalyani, A., Lerner, J., & Tahoun, A. (2021, July). *The diffusion of disruptive technologies* (Working Paper No. 28999). National Bureau of Economic Research. Retrieved from http://www.nber.org/papers/w28999 doi: 10.3386/w28999

Bonfiglioli, A., Crinò, R., Gancia, G., & Papadakis, I. (2024, 11). Artificial intelligence and jobs: evidence from us commuting zones*. *Economic Policy*, eiae059. doi: 10.1093/epolic/eiae059

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018, May). What can machines learn, and what

does it mean for occupations and the economy? *AEA Papers and Proceedings*, *108*, 43-47. doi: 10.1257/pandp.20181019

Callaway, B., & Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, *225*(2), 200-230. Retrieved from https://www.sciencedirect.com/science/article/pii/S0304407620303948 (Themed Issue: Treatment Effect 1) doi: https://doi.org/10.1016/j.jeconom.2020.12.001

Caprile, M., Dente, G., Palmén, R., & Sanz, P. (2015). *Encouraging stem studies – labour market situation and comparison of practices targeted at young people in different member states*. Publications Office. doi: doi/10.2861/519030

Cronert, A. (2019, October). *Is regulatory compliance by employers possible without enforcement? Evidence from the Swedish labor market* (Working Paper Series No. 2019:23). IFAU - Institute for Evaluation of Labour Market and Education Policy.

Deming, D. J., & Kahn, L. B. (2018). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, *36*(S1), 337-369.

Deming, D. J., & Noray, K. (2020, 06). Earnings Dynamics, Changing Job Skills, and STEM Careers*. *The Quarterly Journal of Economics*, *135*(4), 1965-2005. doi: 10.1093/qje/qjaa021

Engberg, E., Hellsten, M., Javed, F., Lodefalk, M., Sabolová, R., Schroeder, S., & Tang, A. (2025). Artificial intelligence, hiring and employment: job postings evidence from sweden. *Applied Economics Letters*, *0*(0), 1–6. Retrieved from https://doi.org/10.1080/13504851.2025.2497431 doi: 10.1080/13504851.2025.2497431

Felten, E. W., Raj, M., & Seamans, R. (2018, May). A method to link advances in artificial intelligence to occupational abilities. *AEA Papers and Proceedings*, *108*, 54-57. doi: 10.1257/pandp.20181021

Firpo, T., Niemann, L., & Danilov, A. (2025, Nov). *The elusive returns to ai skills: Evidence from a field experiment* (Rationality and Competition Discussion Paper Series No. 552).

CRC TRR 190 Rationality and Competition. Retrieved from https://ideas.repec.org/p/rco/dpaper/552.html doi: None

Goldfarb, A., Taska, B., & Teodoridis, F. (2023). Could machine learning be a general purpose technology? a comparison of emerging technologies using data from online job postings. *Research Policy*, *52*(1), 104653. Retrieved from https://www.sciencedirect.com/science/article/pii/S0048733322001743 doi: https://doi.org/10.1016/j.respol.2022.104653

Government Offices of Sweden. (2026). *Sweden's AI strategy* (Tech. Rep.). Government Offices of Sweden. Retrieved 2026-05-21, from https://www.government.se/information-material/2026/02/swedens-ai-strategy/

Green, A., & Lamby, L. (2025, June). The characteristics of the artificial intelligence workforce across OECD countries. *The Indian Journal of Labour Economics*, *68*(2), 541–568. Retrieved from https://doi.org/10.1007/s41027-024-00549-7 doi: 10.1007/s41027-024-00549-7

Grönqvist, E., & Lindqvist, E. (2016). The making of a manager: Evidence from military officer training. *Journal of Labor Economics*, *34*(4), 869–898. Retrieved 2025-12-11, from https://www.jstor.org/stable/26553229

Hellsten, M., Khanna, S., Lodefalk, M., & Yakymovych, Y. (2025). *The effects of artificial intelligence on jobs: Evidence from an AI subsidy program* (IZA Discussion Paper No. 18267). Institute of Labor Economics (IZA). Retrieved from https://ssrn.com/abstract=5774571 doi: 10.2139/ssrn.5774571

Humlum, A., & Bjørnsson Meyer, B. (2022). *Artificial intelligence and college majors* (Study Paper No. 172). The ROCKWOOL Foundation Research Unit.

Håkanson, C., Lindqvist, E., & Vlachos, J. (2020). Firms and skills: The evolution of worker sorting. *Journal of Human Resources*. Retrieved from https://jhr.uwpress.org/content/early/2020/09/16/jhr.56.2.0517-8801R2 doi: 10.3368/jhr.56.2.0517-8801R2

Jensen, M. F. (2023). *Gender differences in returns to skills evidence from job vacancy data and matched employer-employee data* (Working Paper). Retrieved from https://fjaellegaard.com/files/gender_differences.pdf

Mundlak, Y. (1978). On the pooling of time series and cross section data. *Econometrica*, *46*(1), 69–85. Retrieved 2026-04-02, from http://www.jstor.org/stable/1913646

Nybom, M., Plug, E., van der Klaauw, B., & Ziegler, L. (2023). *Skills, parental sorting, and child inequaliy* (Working Paper No. 2023/12). The Institute for Evaluation of Labour Market and Education Policy (IFAU).

OECD. (2023). *Oecd employment outlook 2023*. Retrieved from https://www.oecd-ilibrary.org/content/publication/08785bba-en doi: https://doi.org/https://doi.org/10.1787/08785bba-en

OECD. (2024). *Training supply for the green and ai transitions: Equipping workers with the right skills*. Paris: OECD Publishing. Retrieved from https://doi.org/10.1787/7600d16d-en doi: 10.1787/7600d16d-en

Pouliakas, K., Santangelo, G., & Dupire, P. (2025, December). Are artificial intelligence skills a reward or a gamble? Deconstructing the AI wage premium in Europe. *Eurasian Business Review*, *15*(4), 1091–1128. Retrieved from https://doi.org/10.1007/s40821-025-00302-0 doi: 10.1007/s40821-025-00302-0

Rönnlund, M., Carlstedt, B., Blomstedt, Y., Nilsson, L.-G., & Weinehall, L. (2013). Secular trends in cognitive test performance: Swedish conscript data 1970–1993. *Intelligence*, *41*(1), 19-24. Retrieved from https://www.sciencedirect.com/science/article/pii/S0160289612001055 doi: https://doi.org/10.1016/j.intell.2012.10.001

SCB. (2021). *Arbetskraftsundersökningarna*. Retrieved from https://www.scb.se/publikation/45095

SCB. (2023a). *Ai-användning i företag och offentlig sektor*. Retrieved from https://www.scb.se/publikation/51280

SCB. (2023b). *Arbetskraftsundersökningarna*. Retrieved from https://www.scb.se/

publikation/51637

SCB. (nd). *Arbetskraftsundersökningarna.* Retrieved from https://www.scb.se/hitta-statistik/statistik-efter-amne/arbetsmarknad/arbetskraftsundersokningar/arbetskraftsundersokningarna-aku/

Squicciarini, M., & Nachtigall, H. (2021). *Demand for AI skills in jobs: Evidence from online job postings* (OECD Science, Technology and Industry Working Papers No. 2021/03). Paris: OECD Publishing. Retrieved from https://doi.org/10.1787/3ed32d94-en doi: 10.1787/3ed32d94-en

Tambe, P. (2014, June). Big data investment, skills, and firm value. *Management Science*, *60*(6), 1452-1469. Retrieved from https://ideas.repec.org/a/inm/ormnsc/v60y2014i6p1452-1469.html doi: 10.1287/mnsc.2014.1899

Uddannelses- og Forskningsministeriet. (nd). *Videregående stem-uddannelser.* https://datavarehus.ufm.dk/rapporter/stem. (Accessed: 2024-01-30)

Webb, M. (2020). *The impact of artificial intelligence on the labor market* (Unpublished manuscript). Retrieved from https://ssrn.com/abstract=3482150

Appendix

A Tables

Table A1
AI firm descriptives

	AI firms	Other firms
No. firms	960	22,146
Avg. yearly AI hires	1.6	-
Avg. yearly total hires	33.3	4.9
Avg. No. employees	884.9	130.2

Notes: Notes: This table shows descriptive statistics at the firm-year level of firms with at least one AI hire, and firms with no AI hires. Sample restrictions are imposed.

TABLE A2
Descriptive statistics, hire-vacancy matches

	Obs.	Mean	Median	Std.	Freq., = 1
Full matched sample:					
Mean AI skill	384,022	0.023	0	0.125	
AI skill = 1	384,022	0.013	0	0.111	4,801
AI skill > 0	384,022	0.062	0	0.241	23,675
AI skill $\geq$ 0.5	384,022	0.019	0	0.138	
Number of matched vacancies	384,022	7.995	2	23.993	
Earnings_{t+1}	384,022	439.891	< mean	569.192	
$\text{Log}(\text{Earnings}_{t+1})$	286,710	5.900	> mean	1.137	
Tenure at hiring firm	384,022	2.824	2	1.857	
Gender (1 = Female)	384,022	0.469	0	0.499	
Age	384,022	33.614	31	11.160	
Years of experience	384,022	10.271	8	8.169	
Number of employees	384,022	2,623.672	780	4,374.243	
Year	384,022	2,019.545	2,020	2.682	
Perfectly matched sample:					
Earnings_{t+1}	164,255	442.125	307.300	563.230	
$\text{Log}(\text{Earnings}_{t+1})$	125,999	5.896	5.985	1.110	
Tenure at hiring firm	164,255	2.971	2	1.946	
Gender (1 = Female)	164,255	0.474	0	0.499	
Age	164,255	34.364	32	11.271	
Years of experience	164,255	11.182	9	8.303	
Number of employees	164,255	2,107.741	475	3,796.375	
Year	164,255	2,019.261	2,020	2.753	

Notes: Notes: This table shows descriptive statistics of key variables in the matched hire-vacancy sample. Earnings is reported as annual earnings in thousands of SEK. An observation equals to one hired individual. t is the year of hire. Medians consisting of few observations replaced by size of median in comparison to mean.

Table A3
Skill effects on successful hiring, full table

	Recruitment success				Recruitment time			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.007	0.001	0.005	0.002	-0.059	-0.021	-0.007	-0.030
	(0.034)	(0.033)	(0.015)	(0.006)	(0.200)	(0.189)	(0.118)	(0.068)
Specialized software	0.024	0.021	0.041**	0.001	-0.083	-0.060	-0.259**	-0.059
	(0.029)	(0.029)	(0.016)	(0.006)	(0.124)	(0.121)	(0.131)	(0.084)
Data	-0.052*	-0.051*	-0.021*	-0.009**	0.364*	0.366*	0.181**	-0.025
	(0.030)	(0.030)	(0.012)	(0.004)	(0.203)	(0.200)	(0.088)	(0.032)
Obs.	1,842,282	1,842,282	1,823,090	1,823,087	530,833	530,833	526,555	526,542
Adj. R^2	0.061	0.063	0.460	0.582	0.023	0.028	0.167	0.264
Fixed effects:								
Year		✓	✓	✓		✓	✓	✓
Occupation				✓				✓
Firm			✓	✓			✓	✓

Notes: This table displays estimates of a regression with the outcome variable being the success of posted vacancies measured both in whether the posting led to a hire, and the duration of the recruitment process, with the explanatory variable being the skill content of the vacancy, including an AI skill dummy (1 = AI probability $\geq$ 50%), data skills, software skills, as well at the Deming and Kahn (2018) (not reported). The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation level. *p<0.1; **p<0.05; ***p<0.01.

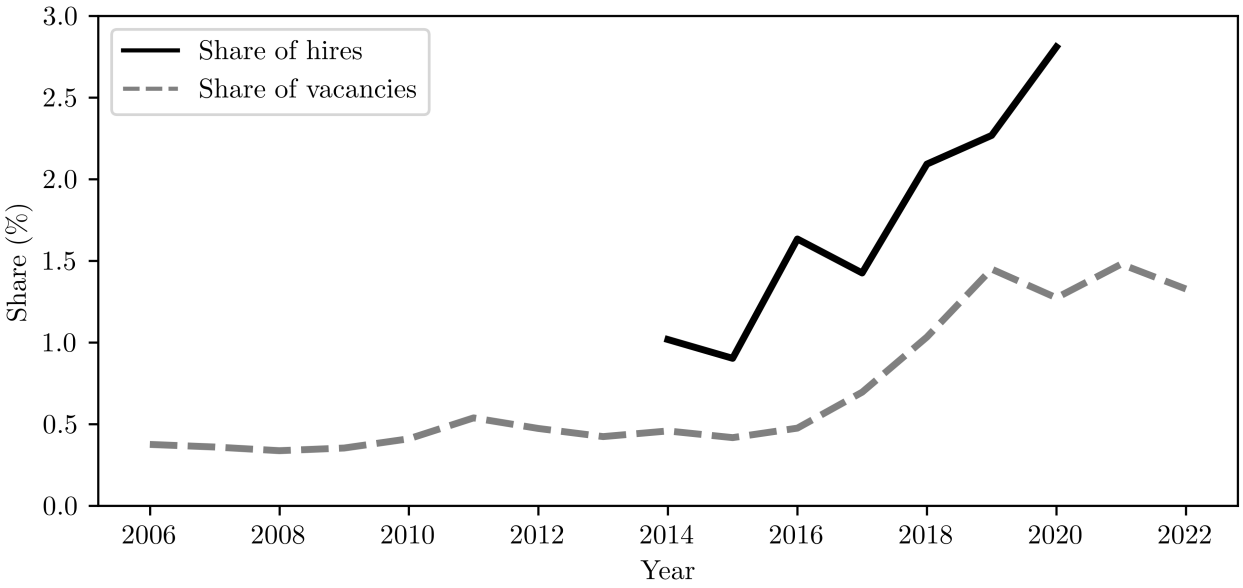
Table A4
Skill effects on successful hiring, AI skills only

	Recruitment success				Recruitment time			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI	0.012	0.003	0.011	0.001	-0.007	0.031	-0.029	-0.034
	(0.042)	(0.040)	(0.019)	(0.006)	(0.255)	(0.242)	(0.120)	(0.066)
Obs.	1,842,282	1,842,282	1,823,090	1,823,087	530,833	530,833	526,555	526,542
Adj. R^2	0.000	0.004	0.451	0.582	-0.000	0.004	0.161	0.264
Fixed effects:								
Year		✓	✓	✓		✓	✓	✓
Occupation				✓				✓
Firm			✓	✓			✓	✓

Notes: This table displays estimates of a regression with the outcome variable being the success of posted vacancies measured both in whether the posting led to a hire, and the duration of the recruitment process, with the explanatory variable being whether the vacancy required AI skills (1 = AI probability $\geq$ 50%). The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation level. *p<0.1; **p<0.05; ***p<0.01.

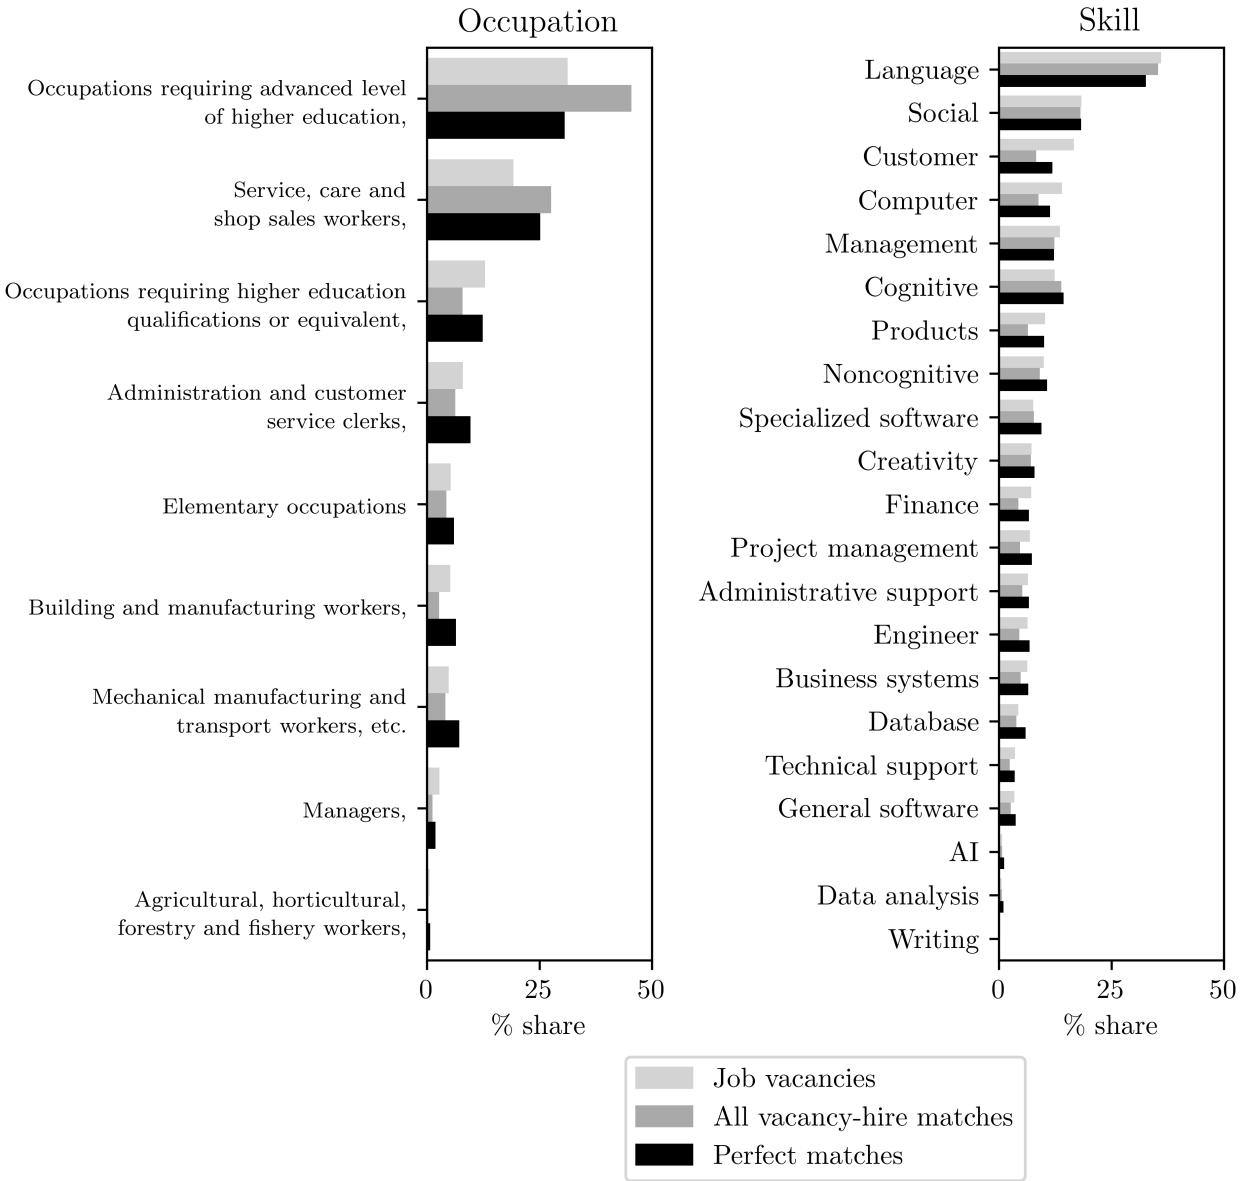
B Figures

Figure A1
AI timeline



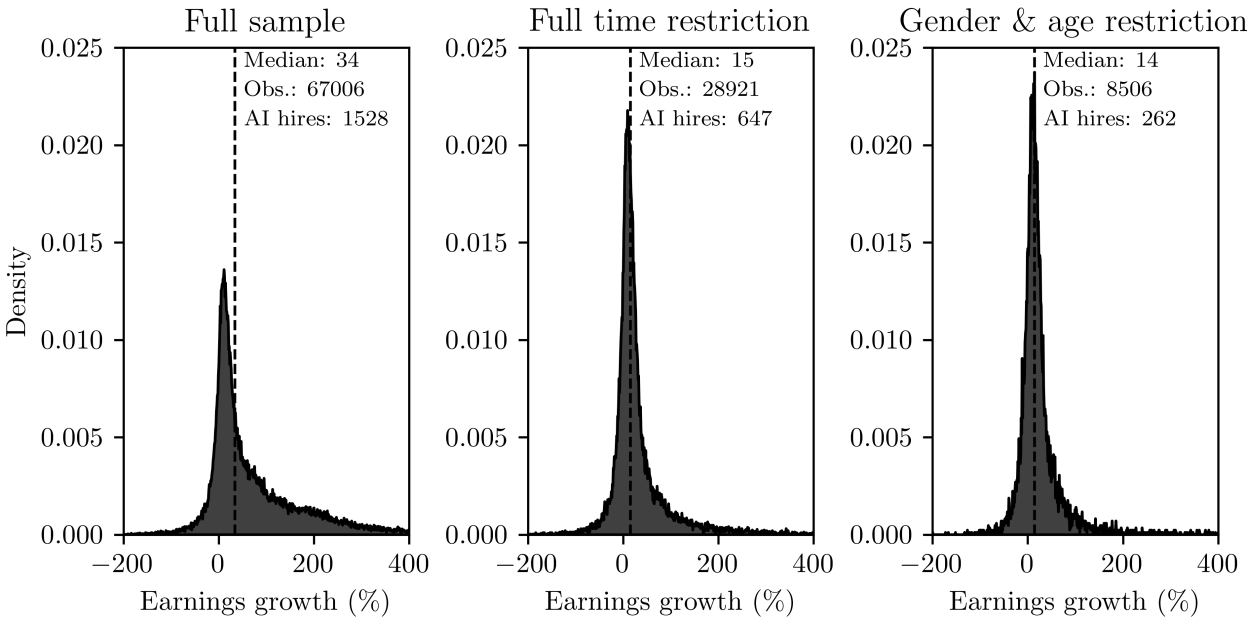
Notes: This figure displays hired AI workers as a share of total matched hires, and AI vacancies as a share of total number of vacancies from workplaces with at least one matched hire.

FIGURE A2
Sample comparison



Notes: This figure is a comparison of differences between the posted vacancies, the vacancies that are in some way matched to a hire, and the sample of perfect vacancy-hire matches. The panel on the left displays occupations, at first level ISCO, and the panel on the right displays skills

FIGURE A3
Earnings growth distributions



Notes: This figure displays the distribution of earninings growth at different sample restrictions. The full sample applies the basic sample restrictions outlined in section 2.2. The full time restriction additionally exluces workers who were not employed for the full calendar year t-1, workers with multiple jobs, and persons who graduated in t or t-1. Gender & age restriction further excludes females and persons outside the age bracket 30-45. Growths of $<$ -100% possible due to log approximation.

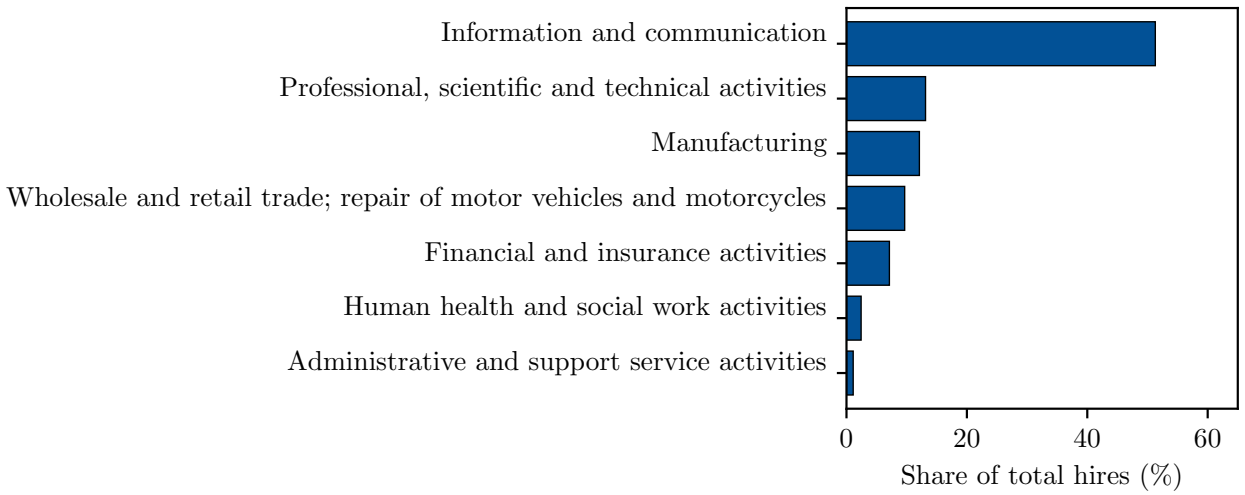
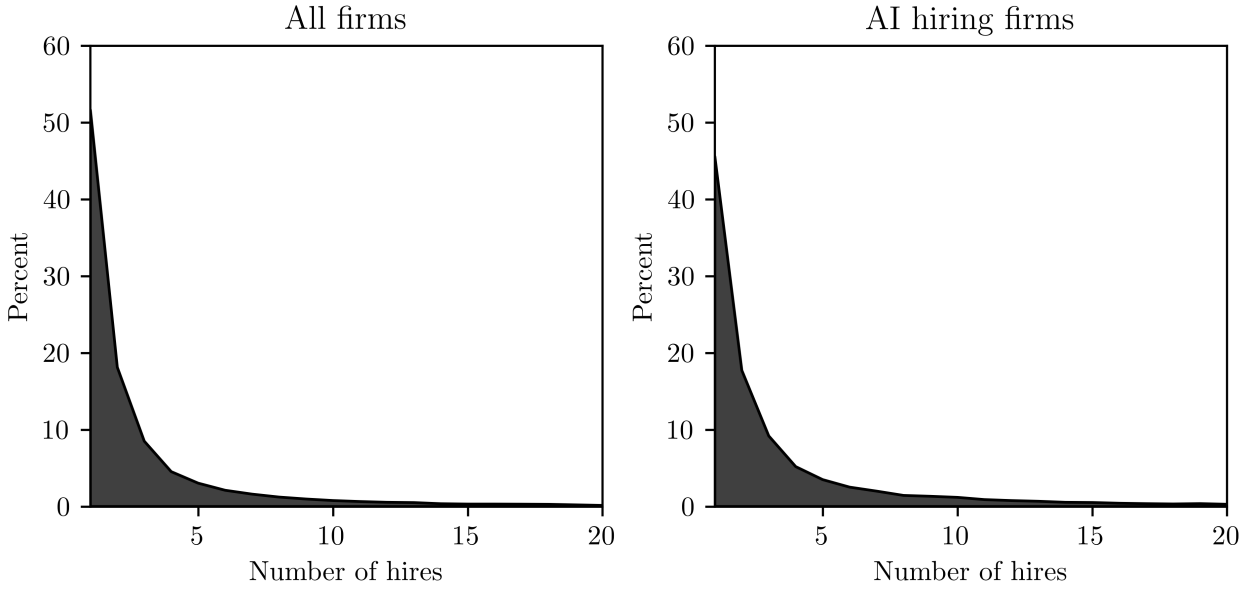


FIGURE A4
AI industries

Notes: This figure displays AI hires by industry, as a share of all AI hires, compared to all hires. Industries at the first level NACE. Industries with less than one percent of all AI workers excluded.

FIGURE A5
Firm number of hires



Notes: This figure displays histograms of the number of hires at the firm level, split into all firms, and firms with at least one AI hire. Observations with 5 or less firms are redacted.

C Sample restrictions

The time period of the matched sample is 2014-2020. The lower bound of this time frame is determined by a change in occupational classification, which would disrupt the matching of hires and vacancies. Data is available up until 2021, which is used to provide the earnings of individuals in 2020, hence the upper limit of the sample.

I make some sample restrictions relating to characteristics of the hires and firms. The objective of these limitations are to filter out workers with a non-representative value of annual earnings, and to make the sample more balanced, and the control group more relevant. The impact of the sample restrictions can be found in Table A5.

Table A5
Sample restriction observations

Step	Number of remaining observations
Matched data 2014-2020	449,391
Minimum one skill	393,852
Ages 19-65	390,256
At least five firm employees	389,412
Private sector only	128,448
Exclude employment agencies	120,510

Notes: This table displays the observations lost at each stage of imposing sample restrictions. Restricting the sample to individuals who stayed in the firm for at least one year is inherent to the matching process, and is therefore not reported.

For individuals, I make the following restrictions: In order to get a representative annual earnings value, I exclude individual who did not stay in the hiring firm for at least one calendar year after the hire. Next, I exclude hires that are matched to vacancies with no specified skill requirements, as these are likely to be vacancies of low quality. I further exclude individuals under the age of 19 and above the age of 65. This decision is grounded in the rarity of full-time work among these age groups, coupled with the limited occurrences of AI-related hires inside these brackets.

I also make the following restrictions at the firm level: To not let any differing connection between control variables and outcome variables of irrelevant industries affect the estima-

tion of control variables, I exclude individuals hired to firms in industries with no AI hires throughout the specified time period. In order not different hiring and compensation patterns of small firms and startups affect the estimates, I also exclude firms with fewer than five employees in the year of hiring. As AI hires are concentrated in larger firms, this does not have a large impact on the number of AI hires in the sample. I also exclude all public sector organizations, defined as entities with municipal or state ownership. This exclusion is motivated by the likelihood of a different nature of AI use in the public sector, as well as wages being more fixed. In fact, some of the public sector occupations where AI skills are most prevalent in Sweden are university occupations such as professors, research assistants, and PhD students (All universities in Sweden are public.), yet these are unlikely to exhibit a earnings premium at all due to fixed wage scales. Finally, I exclude the industry categorized as 'employment activities' (NACE code 78.). This decision is driven by two key considerations: Firstly, it is a common practice in Sweden for companies to outsource their hiring processes to labor hire firms. These firms engage in both the recruitment and subsequent leasing of employees, as well as playing a facilitative role in the hiring process itself. Consequently, the data is biased towards these firms. This is because some vacancies connected to these firms may, in reality, represent hiring activities of other companies. Secondly, labour hire firms often employ temporary workers, meaning the prevalence of part-time work is very high within these firms.

D Hiring date approximation

This section covers the approximation of the hiring date of employees in the years prior to 2019, where only yearly data is available.

Given that it is most common to start a new job at the first day of a month, it suffices to identify the correct month of the hire. I approximate the hiring date by calculating the earnings received by an individual in the starting year as a proportion of their earnings in the second year with the firm. Following this, I then calculate the number of days that

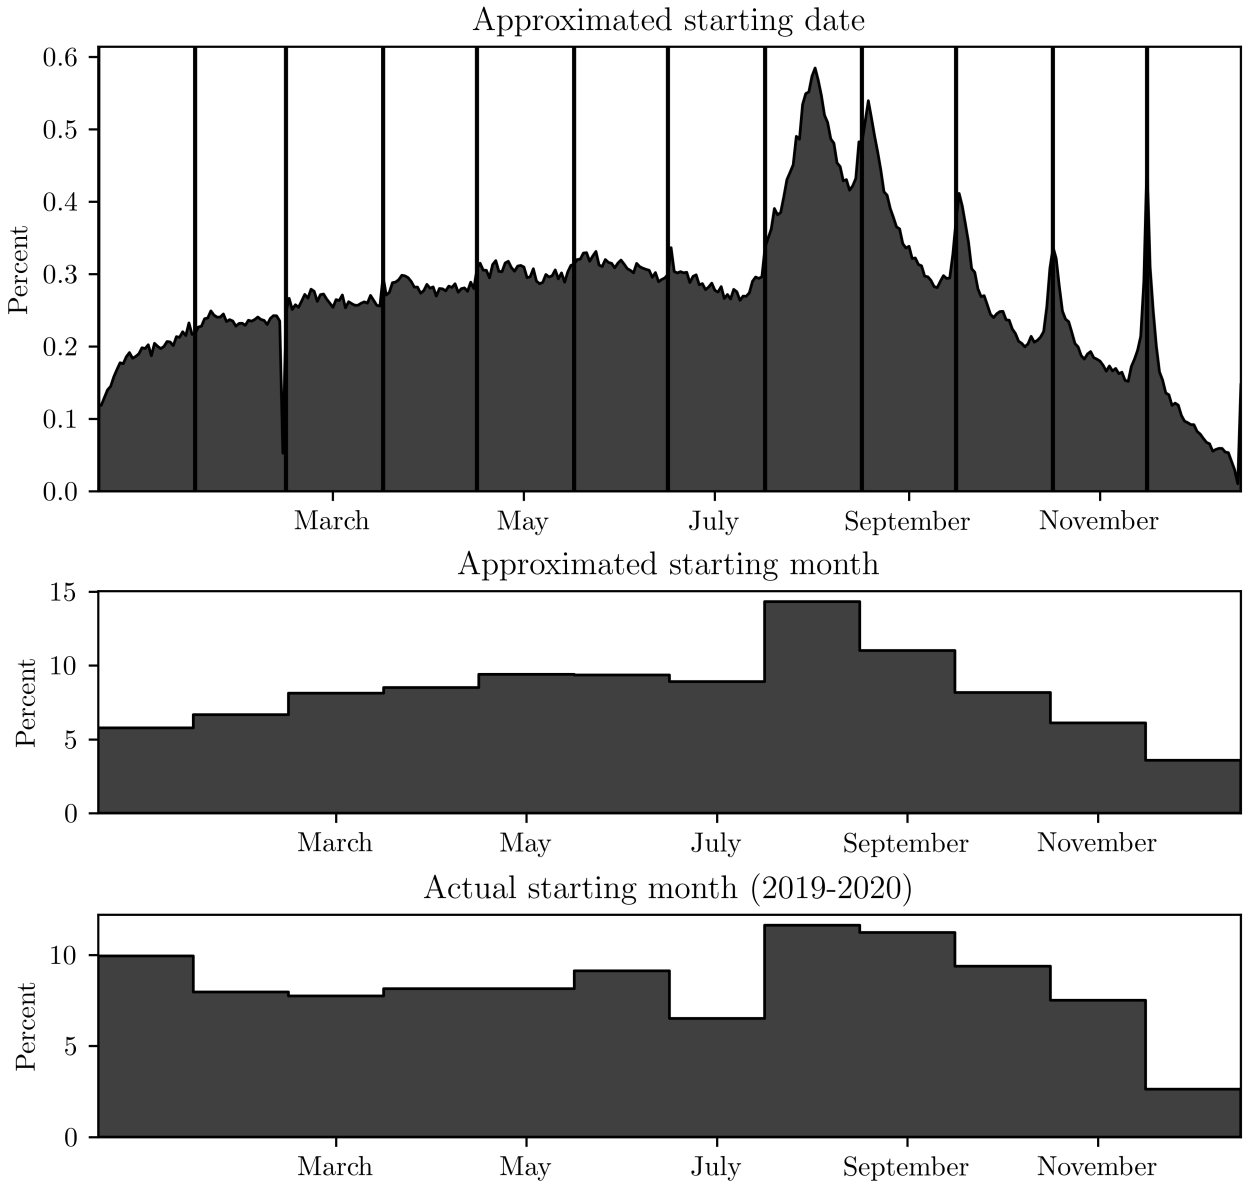
represents their share of a full year's earnings, which subtracted from the full year provides an approximated starting date.

$$\hat{T}_{it} = 365 \left(1 - \frac{E_{ift}^{new}}{E_{it}^{tot}} \right) \tag{5}$$

The rationale underlying this approach is that if an individual earned half their typical annual earnings in the first year, it is reasonable to infer that they initiated their employment roughly halfway through that year.

The graph in Figure A6 provides an overview of the distribution of estimated hiring dates, along with the resulting approximated month, compared to the actual distribution derived from the later monthly data. A well-functioning estimation methodology is expected to exhibit seasonality, particularly towards the commencement of each month, given that such days are commonly preferred as the starting date. While the graph demonstrates discernible seasonality, it appears to deviate slightly from the start of the month in the earlier months. This shifted seasonality can likely be attributed to the approximation not accounting for factors such as wage growth, reduced vacation time in the first year, trial employment periods, among others. Nevertheless, despite these limitations, this methodology serves as a sufficiently robust approach for approximating the starting month.

FIGURE A6
Hiring date approximation



Notes: This figure displays the approximated hiring dates. The top graph is the date approximated by the share of earnings. The second panel turns the approximated date into hiring month, which is used for the final matching to job vacancies. The last panel displays the monthly distribution as recorded in the monthly data available from 2019

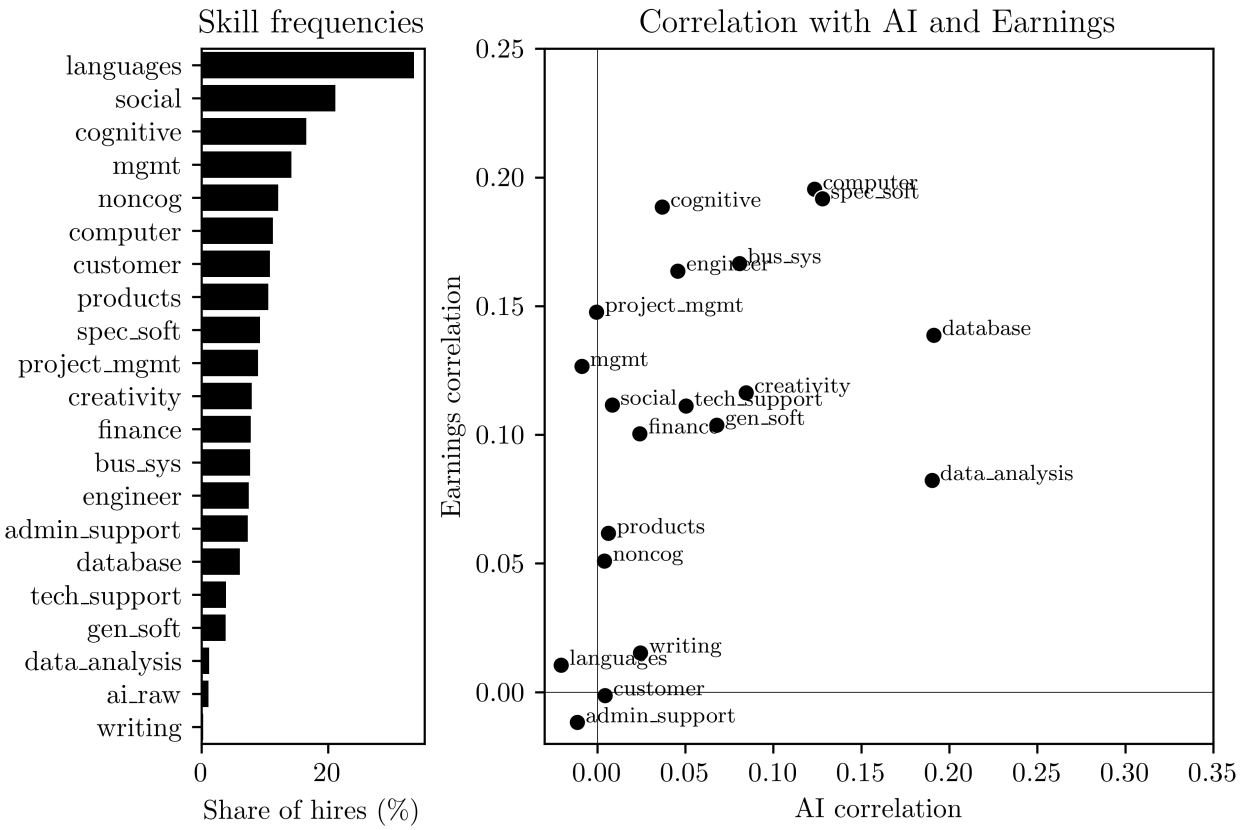
E ADDITIONAL RESULTS

E.1 AI skills or just data and software skills?

AI skills are at the intersection between data and software skills (OECD, 2023). This intersection is substantiated by the high correlation observed in Figure A7, indicating strong

associations between AI, data, and software skills. Because of this, it becomes a challenge to separate the AI earnings premium from that of other correlated skills, especially when all three exhibit moderately high correlations with earnings.

FIGURE A7
Skill correlation



Notes: The panel on the left displays skill frequences in the perfectly matched sample of hires. The panel on the right displays spearman correlation between the Deming and Noray (2020) skill categories and AI skills, as well as with the natural logarithm of earnings.

Table A6 addresses this challenge by separating the effects of AI skills from data and software skills in the same specifications as column (4) and (5) of the baseline regression. The results reveal that AI skills still exhibit a positive effect on earnings when controlling for data and software skills, albeit with a slightly smaller and less significant coefficient. Data and software skills also show positive, but smaller coefficients, aligning with expectations. However, the difference between the coefficients are not great enough to provide evidence on the effects being different as indicated by the non-significant F-tests.

TABLE A6
AI skills, data and software

			Interactions			
	(1)	(2)	(3)	(4)	(5)	(6)
AI skill	0.029**	0.000	0.050***	0.026	0.015	-0.002
	(0.013)	(0.013)	(0.016)	(0.017)	(0.018)	(0.019)
Software	0.015*	0.016**	0.018**	0.018**		
	(0.008)	(0.007)	(0.008)	(0.008)		
Software * AI skill			-0.021	-0.037		
			(0.023)	(0.024)		
Data	0.021**	0.010			0.021**	0.012
	(0.008)	(0.008)			(0.009)	(0.008)
Data * AI skill					0.025	0.004
					(0.025)	(0.026)
F-tests, H_0:						
AI skill = software	0.984	1.148				
P-value	0.321	0.284				
AI skill = Data	0.23	0.397				
P-value	0.631	0.529				
Obs.	76,848	76,848	76,848	76,848	76,848	76,848
Adj. R^2	0.412	0.446	0.412	0.446	0.412	0.446
Control variables:						
Basic individual variables	✓	✓	✓	✓	✓	✓
Firm size	✓	✓	✓	✓	✓	✓
Fixed effects:						
Occupation	✓	✓	✓	✓	✓	✓
Workplace industry	✓	✓	✓	✓	✓	✓
Firm		✓		✓		✓

Notes: This table displays estimates of a regression with the natural logarithm of earnings as the outcome variable, with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. The sample is restricted to individuals who are only matched to a single job vacancy. All regressions use the most restrictive sample. Data combines the skill categories database and data analysis. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

Examining the interactions in columns (3) - (6), the combination of AI skills and software skills yields a negative, albeit non-significant effect. The data interaction seem to affect the results more greatly. The effect of AI skill coefficient disappears, whereas part of its previous magnitude of 4% is shifted over to the interaction with data skills. This observation suggests

a potential disparity in the rewards garnered between AI utilized in the domain of data science and AI developed for broader applications. While the coefficient is not statistically significant, this outcome is unsurprising given the fewer number of observations featuring both data and AI skills, compared to just AI skills.

An obvious concern with this model is distortions due to multicollinearity. A Variance Inflation Factor (VIF) analysis was conducted on Column 1 in order to check this[27]. The findings reveal VIF values of 1.1 for AI skills, 1.66 for software skills, and 1.2 for data skills. These results suggest that multicollinearity is not a substantial concern. Nevertheless, the results of the regression is evidently sensitive to the inclusion or exclusion of skills from the regression model.

Furthermore, it should be noted that the outcomes of this regression analysis may be influenced by the potential oversight of skills that are presumed rather than explicitly stated. This tendency is likely to be prevalent for software skills, as these may be taken for granted in software occupations. It could also conceivably extend to data skills, given the unlikelihood of a position requiring AI skills without either software or data skills.

E.2 Increasing demand for AI skills

As outlined in Section **??**, one possible explanation for the lack of a significant wage competition could be that AI technologies are not yet yielding enough productivity increases to justify engaging in competition for workers with AI skills. This cannot be ruled out, especially considering the sample period ends before the recent advancements in generative AI starting in late 2022. However, there is another period of significant increase in demand for AI skills starting in 2017; see Figure A1 in the Appendix. This demand increase can serve as a proxy for examining the relationship between the earnings premium and spikes in demand for AI skills.

[27] Occupation level transitioned from 4-digit to 2-digit, and industry from the 4-digit to the 1-digit (letter) to facilitate Stata execution.

TABLE A7
AI earnings premium, by period

	(1)	(2)	(3)	(4)	(5)
AI pre-2017	0.132***	-0.003	-0.004	-0.003	-0.027
	(0.020)	(0.020)	(0.023)	(0.023)	(0.027)
AI post-2017	0.231***	0.073***	0.057***	0.056***	0.016
	(0.021)	(0.018)	(0.016)	(0.016)	(0.016)
Obs.	76,848	76,848	76,848	76,848	76,848
Adj. R^2	0.246	0.400	0.411	0.412	0.446
Control variables:					
Basic individual variables	✓	✓	✓	✓	✓
Firm size				✓	
Fixed effects:					
Occupation		✓	✓	✓	✓
Workplace industry			✓	✓	✓
Firm					✓

Notes: This table displays estimates of a regression with the natural logarithm of earnings as the outcome variable, with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). The AI skill dummy is split into hiring before and after 2017. The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. All regressions use the most restrictive sample. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

To investigate this, I set up a regression identical to the baseline regression, but with the AI skill dummy variable divided to indicate whether an individual was hired before or after 2017.

The results display similar but slightly larger coefficients for the period after the demand shock. Prior to the demand shock, there is no evidence of an earnings premium for AI skills when controlling for occupation FE. As anticipated, the results suggest that the demand shock for AI skills appears to drive the observed earnings premium. This implies that the three-fold increase in demand for AI skills is related to the earnings premium increasing from zero to 5%.

E.3 Earnings growth

The preceding sections have highlighted an association between having AI skills and higher earnings, attributing this to the propensity of high-wage firms to recruit workers with AI skills. However, it remains to be answered whether individuals with AI skills experience an earnings increase attributable to possessing these skills, or if they inherently command higher earnings compared to their peers. This distinction is pivotal not only for understanding whether individuals are rewarded and incentivized to transition to firms explicitly seeking their AI skills but also for determining whether firms consider this skill significant enough to warrant higher earnings in order to attract and retain workers possessing this rare skill.

TABLE A8
AI earnings growth premium

	(1)	(2)	(3)	(4)	(5)
AI skill	-0.010	0.016	0.024	0.023	0.015
	(0.026)	(0.026)	(0.025)	(0.025)	(0.024)
Obs.	17,401	17,401	17,401	17,379	17,401
Adj. R^2	0.131	0.146	0.156	0.156	0.176
Control variables:					
Basic individual variables	✓	✓	✓	✓	✓
Firm size				✓	
Fixed effects:					
Occupation		✓	✓	✓	✓
Workplace industry			✓	✓	✓
Firm					✓

Notes: This table displays estimates of a regression with the growth of earnings between the current and last job (approximated by $ln(earnings_{if,t+1}) - ln(earnings_{if,t-1})$) as the outcome variable, with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. All regressions use the most restrictive sample. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

To investigate the earnings growth, I modify the base regression from Table **??**. The outcome variable was adjusted to approximate earnings growth between the last year of the previous

job and the second year of the new job, expressed as $ln(earnings_{if,t+1}) - ln(earnings_{if',t-1})$. As outlined in Section **??**, time-varying controls are shifted to $t-1$, allowing the AI skill variable to retain the effect of matching into higher paying occupations, industries, and firms, thereby indicating whether the individuals gained anything beyond what is expected from changing jobs. The regression uses the same specifications as the baseline regression.

The findings derived from the regression analysis are shown in Table A8. Results do not suggest any association between earnings growth and AI skills.

E.4 AI developers or AI users?

As outlined in Section 1, I categorize individuals possessing AI skills into two distinct categories: developers and users. I make this distinction as I hypothesize that AI developers should experience a higher return to their AI skills than AI users, based on two primary reasons: Firstly, AI developers are tasked with the creation of AI algorithms, a role demanding a high level of knowledge of AI and a substantial investment in skill acquisition. Conversely, AI users predominantly employ AI as a tool within specific tasks. Consequently, it is less probable that firms would be inclined to offer substantial premiums to attract AI users, as the reskilling of existing labor could feasibly accommodate this need. Furthermore, considering that AI users primarily utilize AI for targeted problem-solving, their impact on overall firm productivity should be marginal. Thus, it follows that firms primarily vie for AI developers, given their pivotal role in driving innovation and advancing the technological capabilities of the organization.

In order to separate AI developers from users, I take inspiration from anecdotal evidence that understanding of science, technology, engineering, and mathematics (STEM) seems a requisite for AI development. Notably, job vacancies for AI developers often emphasize qualifications in computer science, mathematics, and even physics. Accordingly, I characterize AI developers as individuals recruited for positions demanding AI skills within STEM

occupations[28] or those possessing a STEM education[29].

TABLE A9
AI earnings premium, STEM interaction

	STEM education			*STEM occupation*		
	(1)	(2)	(3)	(4)	(5)	(6)
AI skill	0.044**	0.034	-0.002	0.044***	0.032**	-0.006
	(0.022)	(0.021)	(0.020)	(0.015)	(0.014)	(0.014)
STEM education	0.004	0.004	0.007			
	(0.006)	(0.006)	(0.006)			
STEM education * AI skill	0.007	0.006	0.012			
	(0.022)	(0.022)	(0.021)			
STEM occupation * AI skill				0.034	0.045	0.081**
				(0.039)	(0.036)	(0.035)
Obs.	76,848	76,848	76,848	76,848	76,848	76,848
Adj. R^2	0.400	0.412	0.446	0.400	0.412	0.446
Control variables:						
Basic individual variables	✓	✓	✓	✓	✓	✓
Firm size		✓			✓	
Fixed effects:						
Occupation	✓	✓	✓	✓	✓	✓
Workplace industry		✓	✓		✓	✓
Firm			✓			✓

Notes: This table displays estimates regressions with the natural logarithm of earnings as the outcome variable, with the explanatory variables being an AI skill dummy (1 = AI probability $\geq$ 50%), and two separate interactions. The left-hand panel includes interactions with the hired person having a STEM education, and the right-hand panel includes an interaction with STEM occupation. The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. All regressions use the most restrictive sample. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

I extend the baseline regression from Table **??** by interacting AI skills with STEM education and STEM occupation, separately. Due to the regression specifications already accounting for occupation, a separate STEM occupation variable is not included, as it would be omitted from the regression. For conciseness, the regression not controlling for occupation FE is omitted from the presentation, as the results are less informative. Similarly, the regres-

[28]STEM occupations are based on the definition of Caprile, Dente, Palmén, and Sanz (2015), following ISCO-08 occupation 21, including associated STEM occupations 31 and 35.

[29]STEM education using the definition of Uddannelses- og Forskningsministeriet (nd), using ISCED 2011 educational subjects 05, 06, and 07, in combination with having a university education.

sion adding workplace industry is excluded as it does not significantly differ from the one introducing firm size FE (columns 1 and 3 from the baseline regression).

The outcomes of the regressions are presented in Table A9. The combination of AI skills and STEM education in columns (1) to (3) do not seem to have a connection to earnings. In fact, STEM education alone does also not display a relation to earnings, with any potential effects being mediated through occupation. Contrastingly, the interaction between STEM occupations and AI skills in columns (4) to (6) is positive, albeit not significant. However, this interaction coefficient increases in size and significance upon the inclusion of firm fixed effects. Interestingly, the coefficient of AI skills alone still falls to zero. This is the only occasion in which I find that individuals with AI skills experience higher earnings than their peers within the same firm. This finding implies that while there is a competitive market for AI skills, it is predominantly concentrated among the limited pool of AI developers. In contrast, the skills necessary for AI utilization do not appear to be as highly valued or sought after in the labor market.

E.5 Recent graduates

The findings thus far suggest that while individuals hired for positions requiring AI skills tend to have higher earnings than others, they do not appear to benefit from entering these jobs. One interpretation of this is that these workers are highly experienced, and thus likely already used AI skills in their previous job, leading to no additional increase specifically tied to entering a new job requiring AI skills. This raises the possibility that AI skills are developed gradually over time and may not constitute an easily acquired ticket to higher earnings. To explore this hypothesis, I investigate whether recent graduates derive any benefits from possessing AI skills upon entering their first job after graduation, as this group of workers could not have used AI skills in their previous job.

I run a regression where AI skills are interacted with a dummy variable for recent graduates. As this analysis aims at capturing the effect on the first job post-graduation, the 'recent

TABLE A10
AI earnings premium, recent graduates

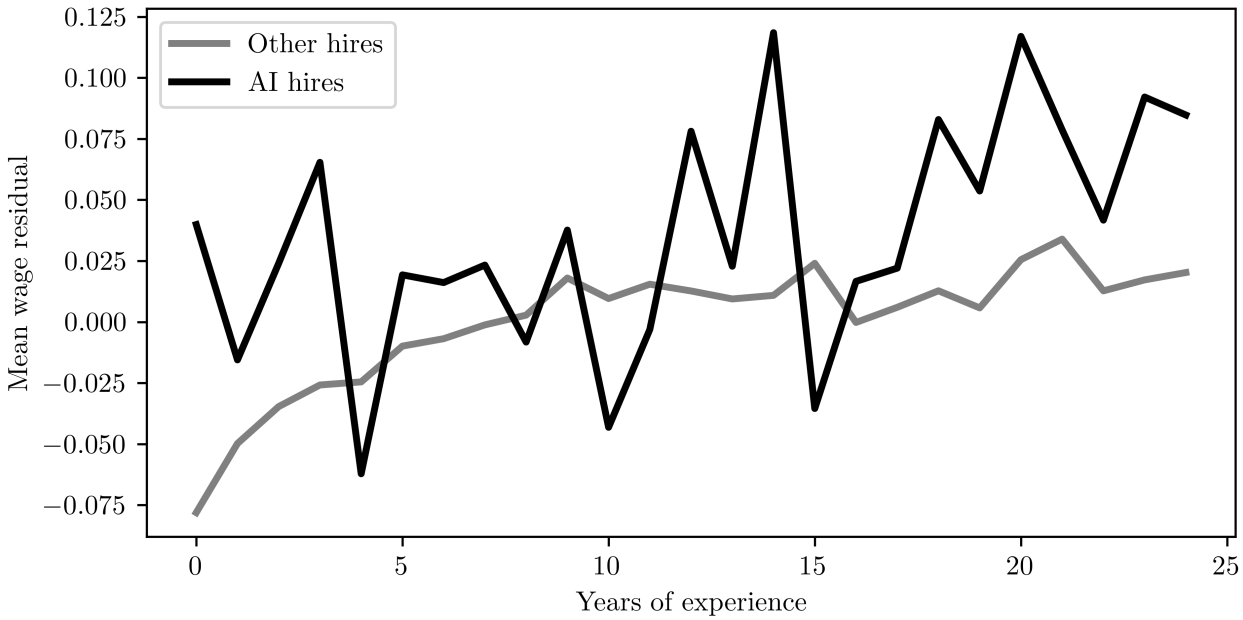
	(1)	(2)	(3)	(4)	(5)
AI skill	0.201***	0.054***	0.046***	0.047***	0.015
	(0.018)	(0.015)	(0.014)	(0.014)	(0.014)
Recent graduate	-0.150***	-0.140***	-0.136***	-0.137***	-0.130***
	(0.007)	(0.006)	(0.006)	(0.006)	(0.007)
Recent graduate * AI skill	-0.016	-0.020	-0.028	-0.030	-0.041*
	(0.027)	(0.025)	(0.024)	(0.024)	(0.023)
Obs.	76,848	76,848	76,848	76,848	76,848
Adj. R^2	0.255	0.408	0.418	0.419	0.452
Control variables:					
Basic individual variables	✓	✓	✓	✓	✓
Firm size				✓	
Fixed effects:					
Occupation		✓	✓	✓	✓
Workplace industry			✓	✓	✓
Firm					✓

Notes: This table displays estimates of a regression with the natural logarithm of earnings as the outcome variable, with the explanatory variables being an AI skill dummy (1 = AI probability $\geq$ 50%), a recently graduated dummy (Individual graduated in t or $t-1$.), and an interaction between AI-skill and recently graduated. The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. All regressions use the most restrictive sample. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

graduate' variable is defined as individuals who graduated in year t or $t-1$.

The results, shown in Table A10, reveal negative but statistically non-significant coefficient estimates associated with the AI skill-recent graduate interaction variable. Taken at face-value. these findings hint that graduates with AI skills benefit less from AI skills than experienced workers, but there is not enough evidence to prove this is the case. The interaction between experience and having AI skills also does also not reveal any evidence in favor of this; See graph of experience and AI skills in Figure A8.

FIGURE A8
Earnings and experience



Notes: This figure displays the relationship between the natural logarithm of annual earnings, and experience, split by hires with and without AI skills. The y axis is the mean residuals of a regression based on specification 4 of the main regression, excluding experience, and experience-squared

F Robustness checks

F.1 AI sample selection

As outlined in section 2.3 the main analysis utilizes a criterion wherein at least 50% of the matched vacancies must possess AI skills for a hire to be classified as having AI skills. In order to make sure the sample does not affect the results, Table A11 presents the regression outcomes when restricting the sample exclusively to perfectly matched hires (i.e., one vacancy per hire). The resulting coefficients are generally similar to the main results, albeit slightly smaller in magnitude. Additionally, the significance levels across the board are somewhat lower.

Recognizing that the chosen cutoff value is somewhat arbitrary, a sensitivity analysis is conducted by iterating through the main regression in specification 4 of Table **??** with varying cutoff values. The coefficients, displayed in Figure A9, reveal that the coefficient remains

TABLE A11
AI earnings premium, perfectly matched sample

	(1)	(2)	(3)	(4)	(5)
AI skill	0.170***	0.037**	0.028	0.027	0.003
	(0.019)	(0.019)	(0.018)	(0.018)	(0.022)
Obs.	33,431	33,431	33,431	33,431	33,431
Adj. R^2	0.264	0.436	0.448	0.449	0.487
Control variables:					
Basic individual variables	✓	✓	✓	✓	✓
Firm size				✓	
Fixed effects:					
Occupation		✓	✓	✓	✓
Workplace industry			✓	✓	✓
Firm					✓

Notes: This table displays estimates of a regression with the natural logarithm of earnings as the outcome variable, with the explanatory variable being an AI skill dummy. The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. The sample is restricted to individuals who are only matched to a single job vacancy. All regressions use the most restrictive sample. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. AKM estimated firm wage represents a control variable based on estimated AKM firm wage fixed effects ($\hat{\psi}_{ft}$). Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

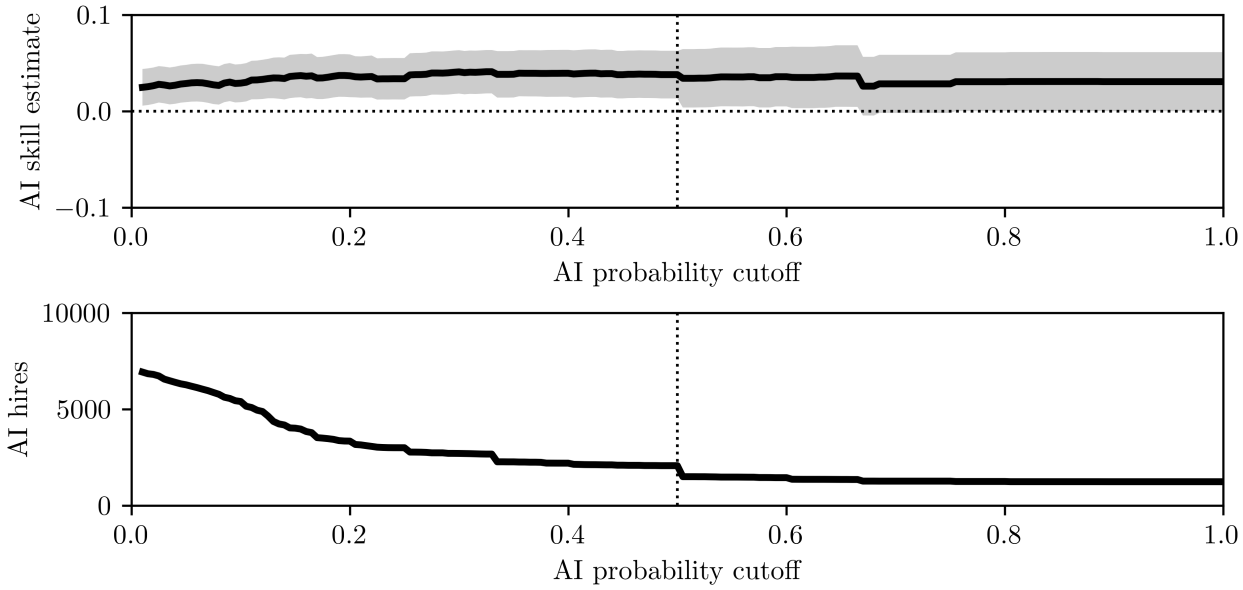
largely consistent and significant, irrespective of the chosen cutoff value.

F.2 Omitted part-time work

An limitation arises from the exclusive use of earnings data, without the availability of hourly wage information. This constraint introduces the possibility that the observed results may be influenced by the prevalence of part-time work. As part-time work is a (mostly) voluntary choice, it would be inappropriate to let this influence the results, as it could lead to bias.

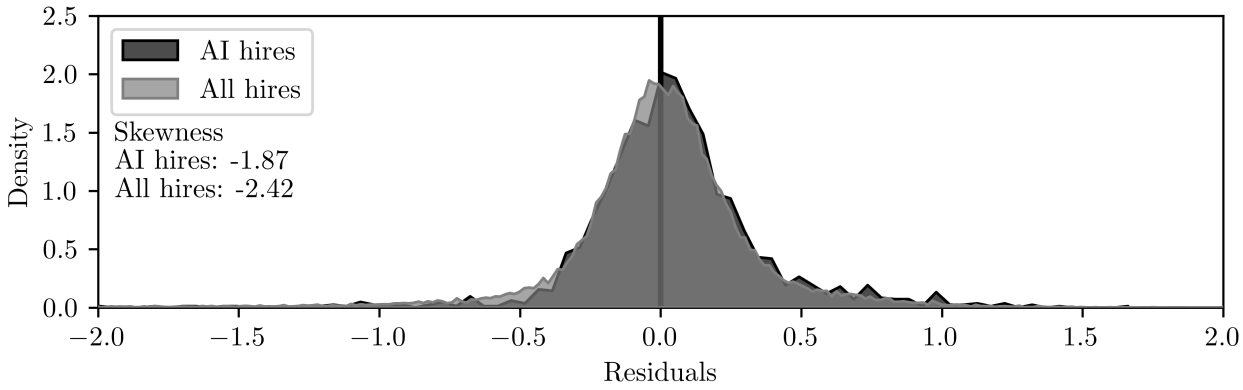
To mitigate this limitation, several steps are taken to minimize the impact of part-time work on the results. First, as outlined in section 2, some sample restrictions are made in order to exclude the majority of part-time workers, I.e. excluding public sector and labour hire firms, as well as individuals with multiple income sources.

FIGURE A9
AI earnings regression, different cutoffs



Notes: This figure displays the estimated effect of AI skill on earnings using different AI probabilities (I.e. mean value of skill of the matched vacancies of the hire.). Based on specification 4 from the baseline regression. Vertical line indicates the AI probility chosen for the main results.

FIGURE A10
Skewness of earnings residuals



Notes: This figure displays the distributions of the residuals of a regression based on specification 2 from the main earnings premium result, with the sample limited to occupations with at least five AI hires. The dark histogram represents the residuals of AI hires, and the light gray histogram displays the residuals of all hires in the sample.

Moreover, the inclusion of control variables gender, age, education, and age of youngest child and the female and age of youngest child interaction is aimed at further address this omitted variable concern. However, despite these efforts, the potential influence of part-time work remains a possible flaw in the model.

To assess whether I fail to control for part-time work, I examine the skewness of the residuals in a regression based on column (2) of the main regression in Table **??**. The choice of specification two is motivated by the inclusion of occupation fixed effects, which is also included in subsequent specifications. As following specifications control for more factors, they are less likely to fail to control for part-time work. The skewness of the residuals could reveal part-time workers, by leaving a left-skewed distribution. This is because the median worker typically works full-time, while a spectrum of part-time work arrangements may contribute to a long left-tail, pulling down the mean earnings.

The observed residuals indeed exhibit a left-skewed distribution, as illustrated in Figure A10. However, the significance of this issue hinges on whether AI hires tend to work part-time less than other workers. To address this concern, I compare hires with AI skills to other hires in the same occupations. The analysis reveals that there is a slight difference in skewness, with AI hires demonstrating less left-skewed residuals. This suggests that workers with AI skills are indeed less prone to working part-time than their peers.

To thoroughly confirm that part-time work does not introduce biases into the results, two robustness checks were conducted. First, workers earning below two standard deviations from the mean earnings within their occupation were excluded, as shown in Table A12. This approach aims to eliminate the lower end of the earnings spectrum, which is more likely to include part-time workers or those with unusually low earnings for other reasons. Second, the sample was limited to male workers in Table A13. This decision is based on the general trend that men are less likely to work part-time compared to women. Excluding the bottom earners did not change the results, whereas the male-only regression left the coefficients positive, although reduced in magnitude, and no longer significant.

F.3 Replicating previous literature

As this paper relates to initial work by Alekseeva et al. (2021), I want to know to what extent the results I find are different due to the different setting (Sweden versus USA), and how

TABLE A12
AI earnings premium, excluding bottom earners

	(1)	(2)	(3)	(4)	(5)
AI skill	0.197***	0.051***	0.040***	0.040***	0.012
	(0.015)	(0.013)	(0.012)	(0.012)	(0.012)
Obs.	76,297	76,297	76,297	76,297	76,297
Adj. R^2	0.269	0.440	0.452	0.453	0.484
Control variables:					
Basic individual variables	✓	✓	✓	✓	✓
Firm size				✓	
Fixed effects:					
Occupation		✓	✓	✓	✓
Workplace industry			✓	✓	✓
Firm					✓

Notes: This table displays estimates of a regression with the natural logarithm of earnings as the outcome variable, with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. Sample excludes hires with earnings under two standard deviations below the mean of the occupation. All regressions use the most restrictive sample. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

much is the differing methodology. In order to investigate this, Table A14 displays the results of replicating the wage premium regressions of Alekseeva et al. (2021) as close as possible, i.e. using only firm and vacancy based control variables. The results are similar to the findings of Alekseeva et al. (2021) when only controlling for firm fixed effects, but similarly to the basic earnings regression in Table **??**, the combination of firm and occupation fixed effects removes the effect of AI skills.

F.4 Unrestricted sample

In the main results, I restrict the sample to that of the specification with the least observations, I.e. the firm fixed effects specification in order to make sure the sample changes does not drive the differences in the results. The results of the unrestricted sample can be found in Table A15. The results do not differ significantly from the baseline results.

TABLE A13
AI earnings premium, males only

	(1)	(2)	(3)	(4)	(5)
AI skill	0.169***	0.036**	0.023	0.023	-0.009
	(0.017)	(0.015)	(0.014)	(0.014)	(0.014)
Obs.	41,956	41,956	41,956	41,956	41,956
Adj. R^2	0.243	0.417	0.433	0.434	0.481
Control variables:					
Basic individual variables	✓	✓	✓	✓	✓
Firm size				✓	
Fixed effects:					
Occupation		✓	✓	✓	✓
Workplace industry			✓	✓	✓
Firm					✓

Notes: This table displays estimates of a regression with the natural logarithm of earnings as the outcome variable, with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. Sample restricted to male hires. All regressions use the most restrictive sample. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

TABLE A14
AI skill regression, without individual traits

	No skills		*Software only*		*All skills*	
	(1)	(2)	(3)	(4)	(5)	(6)
AI skill	0.056***	0.004	0.044***	0.002	0.043**	0.005
	(0.018)	(.)	(0.015)	(.)	(0.018)	(0.007)
Software			0.078***	0.020	0.032*	0.021
			(0.020)	(.)	(0.017)	(0.013)
Obs.	84,032	84,009	84,032	84,009	84,032	84,009
Adj. R^2	0.350	0.420	0.352	0.420	0.353	0.420
Skills					✓	✓
Fixed effects:						
Year	✓	✓	✓	✓	✓	✓
Location	✓	✓	✓	✓	✓	✓
Occupation		✓		✓		✓
Firm	✓	✓	✓	✓	✓	✓

Notes: this table displays estimates of a regression following the preferred specifications of the earnings premia regressions of Alekseeva et al. (2021) with the natural logarithm of earnings as the outcome variable, with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. Location represented by FA-15 municipal groups. Software is the Deming and Noray (2020) specialized software category. Skills denotes the use of a group of control variables based on Deming and Noray (2020), formatted to resemble the Deming and Kahn (2018) skill categories used by Alekseeva et al. (2021) as closely as possible. Categories general software and computer merged to form the computer(general) skill category. Occupation at 4-digit ISCO. Workplace ndustry as third-level NACE. Standard errors in parantheses clustered by location, occupation, and 3-digit workplace industry *p<0.1; **p<0.05; ***p<0.01.

TABLE A15
AI earnings premium, unrestricted sample

	(1)	(2)	(3)	(4)	(5)
AI skill	0.199***	0.048***	0.038***	0.038***	0.005
	(0.015)	(0.014)	(0.013)	(0.013)	(0.013)
Obs.	81,156	81,130	81,122	81,122	76,848
Adj. R^2	0.243	0.396	0.407	0.408	0.446
Control variables:					
Basic individual variables	✓	✓	✓	✓	✓
Firm size				✓	
Fixed effects:					
Occupation		✓	✓	✓	✓
Workplace industry			✓	✓	✓
Firm					✓

Notes: This table displays estimates of a regression with the natural logarithm of earnings as the outcome variable, with the explanatory variable being an AI skill dummy (1 = AI probability $\geq$ 50%). The sample is restricted to hires with at least one skill requirement, and private sector industries (excluding employment agencies, NACE 78) with at least one AI hire. Basic individual variables is a set of individual level control variables in the form of year, gender, experience, experience squared, age, education, location, age of youngest child, and female * age of youngest child. Workplace industry as fourth-level NACE. Occupation at 4-digit ISCO. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.