Returns to AI skills*

Mark Hellsten
mhellsten@econ.au.dk

January 30, 2025
Link to latest version

Abstract

This paper assesses whether workers who develop and apply artificial intelligence (AI) experience an earnings premium. I link skill requirements specified in job vacancies to the individuals ultimately hired to fill those positions using a combination of Swedish job vacancy and matched employer-employee register data. By identifying positions that explicitly necessitate AI skills, this paper seeks to determine if an earnings premium is associated with these skills, while controlling for other individual attributes. Findings suggest a significant earnings premium for individuals hired to positions requiring AI skills. Discerning between AI developers and AI users, the results indicate that the former group experiences a stronger earnings premium. The premium is partly driven by workers being hired into high-wage firms. However, transitioning into roles requiring AI skills does not result in additional earnings increases, indicating that firms do not engage in wage competition for these workers.

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

JEL Codes: D22, J24, M2, O33.

*I am grateful to seminar participants at Aarhus university, University of Chicago Booth School of Business, and University of Chicago, conference participants at AI and the Future of Work (AI at Wharton), Aarhus-Kiel Workshop 2023, and DGPE workshop 2023, as well as Mathias Barding, Anders Humlum, Ina Jäkel, Michael Koch, Magnus Lodefalk, Sarah Schroeder, Bastian Schulz, and Emilie Vestergaard for valuable comments and discussions. Financial support from the Carlsberg Foundation and the Ratio Institute is gratefully acknowledged.

1 Introduction

The current advancements in Artificial Intelligence (AI) technology are anticipated to have transformative effects on the labor market. The skills firms need their employees to possess in order to adopt AI are reportedly scarce (OECD, 2023; SCB, 2023a). This scarcity of a crucial skill, combined with the productivity increases associated with AI, should manifest in higher earnings for the individuals possessing the skill (Acemoglu & Autor, 2011; Acemoglu & Restrepo, 2019). In this paper, I investigate whether firms offer a premium to attract and retain workers with scarce and highly demanded AI skills.[1] The existence and size of such a premium has implications both for the workers involved as well as for the broader adoption of AI, which will be affected by the (wage) costs of investing in this new technology (Foster & Rosenzweig, 2010).

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 that are not possible when relying solely on one of these data sources. First, job vacancy skills at the individual level allows me to directly identify which positions require AI skills, enabling within-occupation comparison of positions with and without AI skill requirements. Furthermore, the methodology allows for an investigation into the driving factors of the earnings premium by analyzing the employment history of hired individuals.[2]

I then determine if AI skills are associated with having higher earnings. I identify a return to

[1]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. I refer to workers with these skills as AI hires and AI workers.

[2]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 the actual recorded earnings filtering out job vacancies that were not successfully filled due to setting the listed wage too low.

AI skills of 5% in earnings, in comparison to individuals hired into similar occupations.

Next, I draw upon Humlum and Bjørnsson Meyer (2022), who find that AI prevalence is related to earnings in college majors concentrated in AI-producing firms, but not in AI using firms. In order to shed light on whether this can be observed at the individual level, I categorize workers with AI skills into two distinct types of workers. Firstly, AI developers encompass dedicated computer scientists who are specifically focused on the development of AI technologies. This group of workers is particularly interesting as they have a stronger connection to the advancement and uptake of the technology within firms. Secondly, AI users consist of professionals in other fields who possess a working knowledge of how to apply pre-designed AI algorithms in their domain. I find some evidence in line with previous literature, indicating a higher return to AI skills for AI developers.

Next, by looking at the wage growth between the current and previous job, I assess whether individuals actually benefit from transitioning to positions where AI skills are required. I find that, despite their higher level of earnings, the individuals hired to fill the positions requiring AI skills do not actually see an additional earnings increase upon entering these jobs. This implies that the higher earnings of these workers carried over from past employment. This result is important, as it suggests that firms do not engage in competition for these workers via attractive wages.

Subsequently, I explore the mechanisms behind the results. The initial results hint at firm sorting playing a role in driving the results, but an analysis specifically investigating this reveals only weak evidence in support of a relation between AI skills and being hired into high-wage firms.

Next, I investigate the reported scarcity of AI skills using the success of vacancy postings requiring AI skills as the outcome variable. I find no evidence in support of this scarcity, which may be an important clue in trying to explain the observed lack of wage competition.

Following this, I shed light on whether the lack of earnings growth is related to workers already exploiting AI skills in their previous jobs. I look at whether new graduates can get a jump start to their earnings by learning a cutting edge skill in their studies, or whether experience is needed to fully develop and benefit from the skill. The results do not provide conclusive evidence on the connection between recent graduation and AI skills.

Finally, I assess whether the lack of wage competition may stem from AI technologies not yet being at a level where the productivity gains outweigh the cost of hiring expensive workers. I analyze the premium before and after a period of increased demand for AI skills which started in 2017. The findings suggest that the AI earnings premium is positively related to the demand for AI skills.

A closely related paper to my study is Alekseeva, Azar, Giné, Samila, and Taska (2021), who use US job vacancy data to investigate the demand for, and wage premium of AI skills, revealing an associated wage premium of 5% compared to vacancies posted within the same firm and occupation. (11% without occupation fixed effects).

In terms of methodology, the most closely related paper to this is Jensen (2023), who matches job vacancy skills to register data in order to uncover gender-based differences in skill wage premiums. Whereas the matching of job vacancy data to register data has been done before at the firm-level,[3] Jensen (2023) is, as far as I am aware, the only other paper conducting matching between skills from job vacancies and hires at the individual level. As the focus of my study is a rather uncommon skill, as well as substantial differences in the data, I develop the matching approach independently instead of directly replicating the method of Jensen (2023). However, some inspiration is drawn from his paper, particularly in forming the econometric approach, given the similar nature of the final data.

My paper can be linked to 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

[3] E.g. Bagger, Fontaine, Galenianos, and Trapeznikova (2022)

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[4], suggests that AI skills are utilized to complete work tasks created by new technologies. This underscores the rationale behind anticipating an earnings premium, as new tasks are more complex and necessitate higher skill levels for their completion.

There is also a host of related literature investigating methods to link advancements in AI technology to labor (Brynjolfsson, Mitchell, & Rock, 2018; Felten, Raj, & Seamans, 2018; Webb, 2020), and AI effects on firm outcomes (Babina, Fedyk, He, & Hodson, 2024). My paper is also related to Acemoglu, Autor, Hazell, and Restrepo (2022), who investigate the effects of AI exposure on labor, finding that firms exposed to AI decrease their hiring, as well as change their demand for skills. Similarly, Bonfiglioli, Crinò, Gancia, and Papadakis (2024) find that AI adoption only leads to positive labor market outcomes for high paid workers and STEM occupations.

Finally, I also draw upon Alekseeva, Gine, Samila, and Taska (2020), Baruffaldi et al. (2020) and Deming and Noray (2020) to create a compilation of AI skill keywords, as well as skill categories of Deming and Kahn (2018) and Deming and Noray (2020) to benchmark AI skills against.

The paper is outlined as follows: In Section 2 I discuss the data sources and sample limitations, along with the definition of AI skills, the methodology of matching job vacancies to individuals, and some characteristics of AI hires. After this, I outline the econometric model in Section 3. Section 4 presents the baseline earnings premium regression results, which I investigate the mechanisms behind in Section 5. I conclude with Section 6.

[4] As seen in Figure A1 in the Appendix

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. One common application is the use of AI-powered tools developed to assist workers (Lane & Williams, 2023). Examples of these tools include generative AI and predictive algorithms. However, in scenarios where an individual uses AI as a tool without needing to understand any underlying mechanisms, such a worker is not classified as an AI worker within the context of my study. Consequently, the tasks performed by these individuals are not considered to require AI skills as defined here.

In order to access skills required for positions, I utilize online job vacancy data sourced from the Swedish Public Employment Service, a predominant platform for job postings in Sweden. 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 et al., 2020; Baruffaldi et al., 2020; Deming & Noray, 2020).[5] 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. 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. From this, I have a list of skills for each job vacancy. I then use text matching against the list of AI keywords to create an indicator variable representing whether a job vacancy requires AI skills.

2.2 Register data and sample

My study also incorporates register data that encompasses the entire population of individuals and firms in Sweden. The register data provides earnings of hired individuals.[6] Moreover, this data serves as the primary source for individual and firm-level control variables. Additionally, the panel structure of the data, covering the years 1990-2021, enables the tracking of individuals over time. This allows me to find out more about who workers with AI skills are, as well as using their previous job for comparison.

The sample of matched data spans from 2014 to 2020. As the data is being matched to job vacancies, I only include new external hires. Thus, incumbent workers as well as internal

[5]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.

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

hires are not considered. Some sample restrictions are imposed with two primary objectives: First, I make restrictions chosen to give the 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 less than five employees, and to hires in industries that had at least one AI hire in the full time period. Second, as the primary outcome variable is earnings, I do what is possible to 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-time workers, I exclude workers who were employed in more than one job simultaneously.[7] Finally, as the skills of an individual are less likely to affect earnings in public sector, I only consider hires in private 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.

The process of matching vacancies with hires is conducted through the utilization of four key identifiers: the firm, occupation, location, and the time period of both the job vacancy and the hire. This matching process presents some complexities, leading to two main challenges:

First, while the date of the job vacancy publication is known, the precise timing of the subsequent hiring (if it occurs) remains undetermined. Secondly, a singular firm may advertise

[7]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.

multiple matching vacancies in a given time period, thereby generating numerous potential matches for a single hire. The subsequent sections detail the methodology employed in the matching process. This methodology is crafted with emphasis on alleviating the aforementioned issues.

In order to match job vacancies to hiring, I need to obtain the start date of new hires. 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. Details on this hiring date approximation can be found in Section D in the Appendix.

As job advertisements typically do not explicitly mention the intended starting date of the position, the remaining piece in the time frame puzzle pertains to determining the duration before the hiring month within which to consider posted vacancies. To establish a reasonable time frame for associating vacancies with hiring, a regression is conducted at the firm-location-occupation-month level. In this analysis, the dependent variable is the occurrence of a hire, and the explanatory variables comprise indicators for relevant job vacancies being posted in each month leading up to the hire. The graph in Figure 1 displays the result of the regression. In line with a similar exercise on Danish data by Bagger et al. (2022), my results indicate that hiring is most strongly associated with vacancy postings within the five months preceding the hiring month. Consequently, for the purpose of matching hires to vacancies, I consider vacancies posted up to five months prior to the hiring month.

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 five-month window preceding the hire.

Given the focus on a relatively rare skill, a rigorous approach is adopted to ensure the precision of the vacancy-to-hire matching. This involves an iterative 'filling the vacancies'

Figure 1
Hiring period regression

Notes: This figure displays the coefficients of a regression where an observations represents a cell consists of one firm, occupation, location and month. The dependent variable is a dummy variable indicating whether there was a hire in the cell, and the explanatory variables displayed on the x-axis are dummies representing vacancy posting in the same cell for each of the 13 months leading up to the hire (with zero being the month of hiring.). The regression controls for month fixed effects.

process. Initially, if a single hire is connected to a single vacancy, that pairing is confirmed, and both the individual and the vacancy are excluded from the process. The process then continues, revealing more one-to-one matches as a result of the previous round of matches having been allocated and removed from the process. This iteration is repeated until no additional matches are feasible.[8]

Following this, I add an additional matching process, where I allow a hire to be matched to multiple vacancies. In these instances, the hire is assigned an average value of the skill content of the associated job vacancies. This third step, while lower in quality, is necessitated by the fact that larger volumes of hiring and vacancy posting will lead to lower number of clean matches, meaning that the matching process does not work as well for large firms. The

[8]This process is run using 4-digit ISCO occupations. Subsequently, the matching criteria are relaxed to a broader occupational classification, shifting from 4-digit ISCO to the more aggregated level of 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. This second step is done in iterations in the same way as the first, stopping when there are no more additions, and as many clean matches as possible have been established.

hiring of large firms is important to consider, as workers with AI skills are concentrated in these firms.

However, there are ways to mitigate the lower quality of the one-to-many matching. The average value of the indicator variables corresponding to skill categories created in the last matching step can be interpreted as the probability that a specific skill was required for a hire. Utilizing this approach allows for the establishment of a threshold probability for determining whether a hire possesses a certain skill. For all skills, including AI skills, the cutoff is set at a minimum probability of 0.5. This implies that hires are classified as having AI skills if they either perfectly match a vacancy requiring AI skills or if they are associated with a mix of vacancies, where at least half of these vacancies necessitate AI skills.

The number of matched hires and vacancies are displayed in Table 1. About 6% 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). Vacancies with AI skill requirements were successfully tied to hires in 28% of cases.

Table 1
Hire-vacancy matching numbers

	Sample with AI skills	Full sample
Hires	-	1,379,591
Vacancies	28,232	2,628,961
Matched vacancy-hires	6,778	88,440
Hires with AI skill probability $\geq 50\%$	2,045	-
Perfect matches	1,061	40,156

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

It is possible to use the perfectly matched sample for robustness checks. This is warranted by some disparities in skill content and occupational distribution between the full sample and perfectly matched sample.[9] There is an over representation of high-skill occupations in the full sample. This pattern can likely also be attributed to such occupations typically being situated in larger firms, from which fewer workers are perfectly matched to vacancies due to the higher likelihood of these firms posting multiple similar vacancies within the same time period.

2.4 Workers with AI skills

The matched data consists of hires, of whom some are attributed AI skills. Table 2 presents some basic characteristics of the AI hires.[10] 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 control variables, it is observed that AI workers typically have higher earnings compared to their non-AI counterparts, with the mean being about 35% 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 2
Mean AI worker characteristics

	AI hires	Other hires
Earnings, $t+1$	572.062	422.688
Age	34.452	34.722
Experience	10.703	10.877
Share tertiary education	0.800	0.565
Share female	0.316	0.435

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

This higher education and earnings align with the predominant occupations of AI workers, primarily concentrated in various forms of software development, as detailed in Table

[9]See comparison in Figure A2 in the Appendix

[10]Full descriptives of the different samples used can be found in Table A1 in the Appendix.

2. Most notably, AI workers are highly concentrated within the occupation group of ICT architects[11], although they are found in most occupations where tasks include data analysis or programming.

FIGURE 2
AI occupations

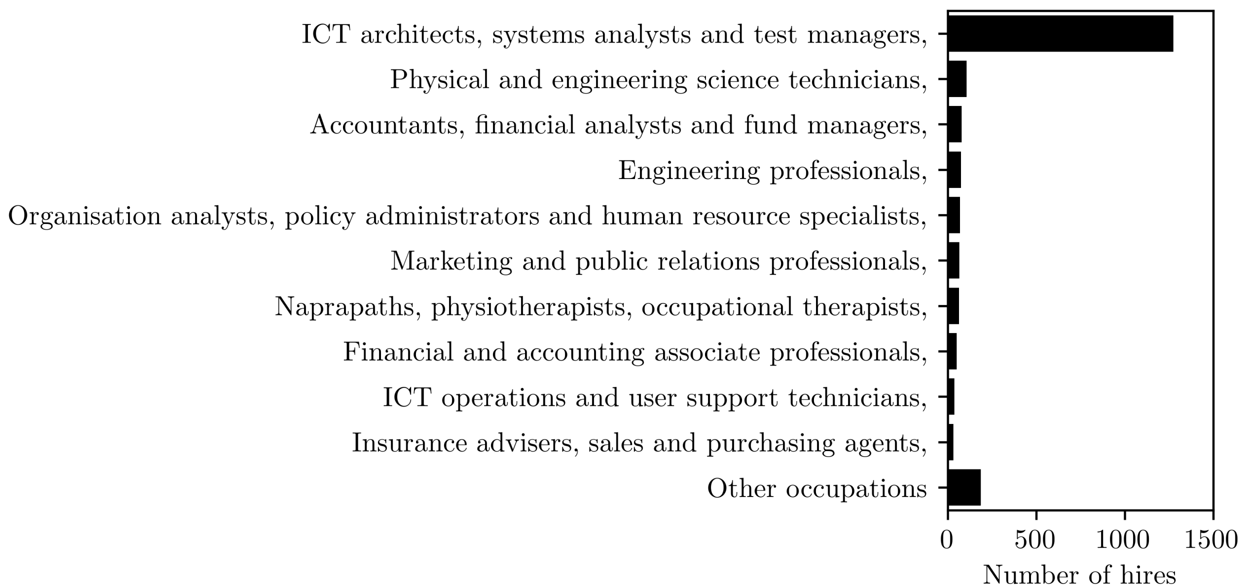

Notes: This figure displays the most common occupations of hires with AI skills. Occupations reported as 3-digit ISCO.

Given this occupational distribution, the most prevalent skills among AI workers, as depicted in Figure 3, 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.

The industry distribution of AI hires[12] shows that AI hires are over-represented in the information and communication sector, as well as the finance and insurance sector. Other than that, AI hires are mostly distributed similarly to the rest of the workforce.

[11] 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.

[12] See Figure A4 in the Appendix

FIGURE 3

AI hire skills

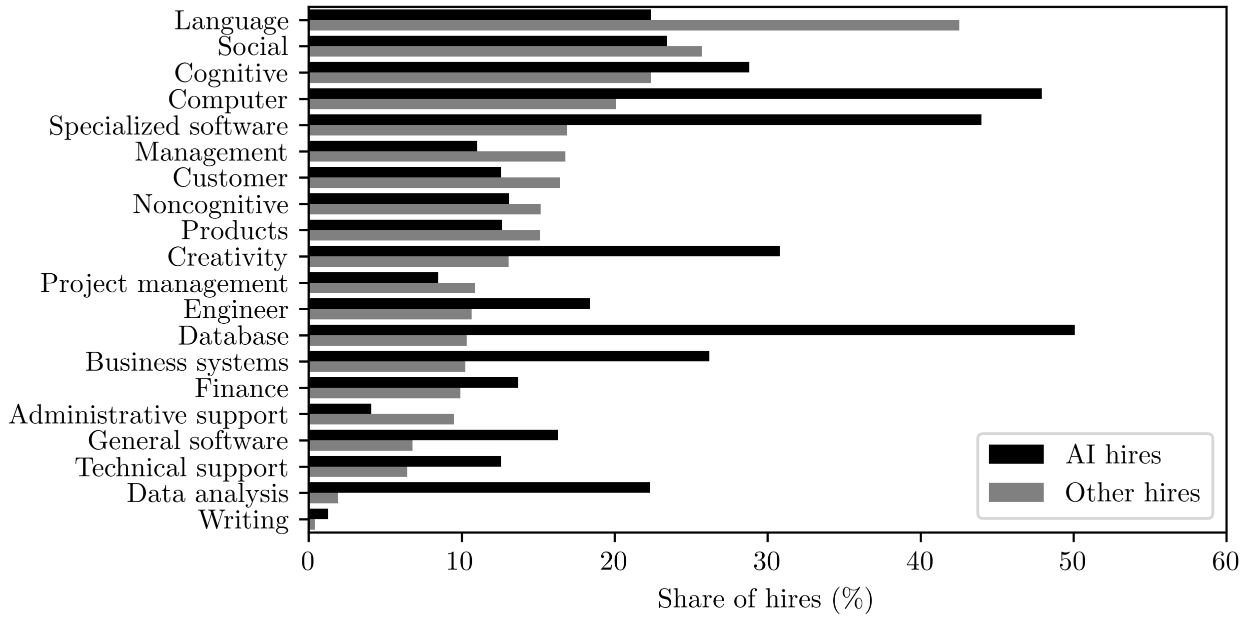

Notes: This figure displays the skill content of perfectly matched vacancies, split by hires with and without AI skills. Sample limited to occupations with at least one AI hire.

TABLE 3

AI firm descriptives

	AI firms	All firms
No. Firms	379	9,125
Avg. No. AI hires	5.257	-
Avg. No. All hires	91.041	9.692
Avg. No. Employees	802.653	136.201

Notes: This table shows descriptive statistics at the firm level of the full sample of firms (with at least one matched hire), and AI firms (Firms with at least one matched AI hire in the period).

In examining the attributes of firms engaging in the recruitment of individuals possessing AI skills, Table 3 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 4,547 distinct firms that have advertised at least one job vacancy related to AI. Within this cohort, 834 firms have successfully hired at least one candidate whose recruitment is, at least partially, associated with an AI job posting. Owing to sample restrictions, there are 379 distinct AI hiring firms in the final sample.

3 Econometric approach

The basis for the empirical analysis is the data of new hires matched to skill content of vacancies. 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 means I want the control group of individuals hired to fill positions requiring AI skills to be other new hires. While this is panel data, as individuals can appear multiple times, the fact that they have to switch jobs within firms in the sample, and successfully be matched to job vacancies multiple times within the 2014-2020 time period in order to do so, most individuals only appear once. Because of this, the data can, in practice, be thought of more like a pooled cross section. The baseline regression uses the model in Equation 1.

$$ln(earnings_{if,t+1}) = \beta_0 + \beta_1 AIskill_{ift} + \boldsymbol{\beta_x X_{it}} + \boldsymbol{\beta_z Z_{ft}} + \epsilon_{ift} \tag{1}$$

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.

Unless otherwise stated, all regressions include control variables for gender, age, work experience, experience squared, and education level. In addition, fixed effects are integrated for year, location[13], age of youngest child, and the interaction of female gender with the age of the youngest child.

There are some additional factors that should possibly be controlled for, but through which an earnings effect of AI skills could be transmitted. In order to enable an overview of how AI skills interact with earnings through these variables, I add these individually in subsequent columns.

[13] FA-15 municipality groups.

Occupation FE are added in the second column. Occupation is perhaps the most crucial control variable, as AI hires tend to be concentrated in high-wage occupations. The specification introducing occupation is therefore the preferred one in terms of estimating individual benefits of AI skills There is, however, a possibility that an individual was hired into a specific occupation solely due to possessing AI skills. Thus, there is some risk that part of an AI skill effect is captured in the addition of occupation, and it should therefore be acknowledged that the true earnings premium may lay somewhere between the specifications including and excluding occupation FE.

AI skills may be more in demand in certain kinds of workplaces and firms. Because of this, I view some firm characteristics as possible transmission mechanisms from AI skills to earnings. These variables are workplace industry, firm size, and firm fixed effects. In order to investigate how these variables affect the results, I add these variables to separate columns.

Of the additional control variables, firm wages are especially likely to play an important role in any earnings premium results. Therefore, the inclusion and interpretation of firm fixed effects become important. A significant result when not including firm fixed effects would be interpreted as there being an earnings premium when compared to the average worker in the economy, speaking towards the individual benefits of AI skills. The same results when including firm FE, on the other hand, should be interpreted as an earnings premium compared to the peers within the same firm, which speaks more towards whether firms value the skill enough to warrant higher pay for the individual employee, as well as whether sorting into high paying firm is an important mechanism. I make sure that there is enough variation in the data for the inclusion of firm FE by looking at the distribution of firm number of hires, it is revealed that while a large share of the firms do have few hires, the majority of individual observations are concentrated within larger firms.[14] Singleton observations within firms are rare, and this rarity is even more pronounced within firms engaged in AI hiring.

[14]See histogram in Figure A5 in the Appendix

This observation suggests that limited within-firm variation in the data should not be a crucial issue.

I cluster the standard errors by occupation*firm*year, following the approach of Jensen (2023), who in turn follow the advised clustering of Hersch (1998) and Cameron and Miller (2015).

In certain regressions, my focus shifts from examining the level of an outcome variable in the year following the hiring, to exploring the change in this variable between the previous job, and the new job. For these specific models, I adjust all time-varying control variables to their values at time $t-1$, i.e. the year or firm prior to the hire. This modification is done in order to preserve potential sorting effects that might be present, as to provide a clearer understanding of the actual impacts experienced by individuals who transition between jobs.

The number of observations is significantly lower in the regressions using earnings growth as the outcome variable. This reduction in sample size can be attributed to two factors. The first factor stems from the incorporation of occupation FE. As mentioned in the section detailing sample restrictions in Section, C in the Appendix, the occupational classification system I employ in this paper was implemented in 2014. Consequently, obtaining data for the control variable for the time period $t-1$ is not possible for individuals hired in 2014, necessitating the exclusion of hires from this year. The second factor comes from the need to exclude all individuals who were not employed in the year preceding the hiring year ($t-1$). This exclusion applied to those who were unemployed during $t-1$, including individuals who were transitioning out of unemployment, as well as recent graduates entering the workforce.

Moreover, as my data comprises earnings data rather than hourly wages, there are some disproportionately high growth numbers influenced by individuals changing their connection to the labor market. Previous research dealing with similar challenges have implemented sample restrictions to confine the analysis to individuals with an active connection to the labor market. For instance, Frederiksen, Halliday, and Koch (2016) confines their sample to individuals aged 35-40, splitting the sample by gender. Bartel and Borjas (1981) employs a sample of men not transitioning into retirement or education for the same purpose. Although I could adopt such broad limitations, I face a trade-off between the number of observations and accuracy, particularly when combined with the previously mentioned restrictions.

As an alternative I impose sample restrictions specifically targeting individuals with a limited connection to the labor force. Initially, recent graduates are excluded, given their likelihood of going from a part-time student job to full-time. Subsequently, similar restrictions are applied to $t-1$ as the ones applied to t in Section 2.2, namely excluding workers with multiple jobs and individuals not employed for the entire calendar year in the departing firm. While there are still instances of workers exhibiting unreasonable growth rates, the disparities in distributions between the alternative sample restrictions are not substantial.[15]

Consequently, all these restrictions lead to a reduction in the number of observations when using growth from the previous job as the outcome variable, particularly due to the large numbers of recent graduates, and hires not employed for a full calendar year in their previous job.

4 Baseline results

Table 4 presents the findings of the baseline AI skill earnings premium regression. The AI skill dummy variable has a positive and statistically significant coefficient across the first four columns. The first column indicates an earnings premium of 22%[16] , which then, as expected, falls significantly with the inclusion of occupation FE, leaving a 5% premium. The further addition of industry FE in the third column results in an additional decrease in the coefficient size, although only by one percentage point. This suggests that a portion of the observed earnings premium can be attributed to employment within industries that typically offer higher wages. On the other hand, the introduction of firm size as a control in the fourth

[15] See histograms of earnings growth of the different samples in Figure A3 in the Appendix.

[16] Earnings premium calculated as $(e^{\hat{\beta}_1} - 1) * 100$ in regressions using $ln(earnings_{if,t+1})$ as outcome variable.

column does not change the AI skill coefficient at all.

TABLE 4
AI earnings premium

	(1)	(2)	(3)	(4)	(5)
AI skill	0.199***	0.048***	0.038***	0.038***	0.005
	(0.016)	(0.014)	(0.013)	(0.013)	(0.013)
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 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.

In column five, firm fixed effects are introduced, resulting in the effect of the AI skill dummy variable turning insignificant. The interpretation of this is that the earnings premium observed in previous columns is primarily related to individuals with AI skills are being hired by high-paying firms. I revisit the implications of this result in section 5.1.

4.1 Robustness checks

Robustness checks can be found in Section E in the Appendix. The chosen matching method of using a 50% cutoff for AI skills is scrutinized in Section E.1, wherein I find that the results hold when using the perfectly matched sample, as well as any other cutoff value. Subsequently, in Section E.2, I evaluate the potential impact of not controlling for part-time employment, finding similar coefficients. In section E.3, 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 E.4, 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 E.5, 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.

4.2 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 exploit the 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[17] or those possessing a STEM education[18].

TABLE 5
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 4 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

[17]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.

[18]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.

is omitted from the presentation, as the results are less informative. Similarly, the regression 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 5. 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.

4.3 Wage competition

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.

To investigate the earnings growth, I modify the base regression from Table 4. The outcome variable was adjusted to approximate earnings growth between the last year of the previous

TABLE 6
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.

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 3, 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 6. Results do not suggest any association between earnings growth and AI skills[19]. To make sure this result is

[19] I also run this growth regression split by AI developers and users, as in Section 4.2, as AI developers seems to experience a higher return to AI skills than AI users, in terms of earnings level. Consistent with the previous earnings growth regression, I shift the STEM variables to period $t-1$. The result of this regression can be found in Table A2 in the Appendix. I find positive but nonsignificant coefficients for AI developers. Similarly to the main growth regression, the results provide no evidence that entering a job requiring AI skills corresponds to an increase in earnings. It should be acknowledged, however, that the growth sample exhibits limited variability in the case of STEM occupations with AI skills. Only 80 instances involve hires with AI skills originating from a STEM occupation, see Appendix, Table A3.

not affected by the extreme earnings growth rates outlined in Section 3, I run a robustness check in Section E.6 in the Appendix. The results are not found to be driven by extreme growth values. The implications of the lack of observed growth are investigated further in the next section.

5 Exploring mechanisms

This section explores the mechanisms driving the results observed in Section 4. I first investigate whether workers entering high paying firms plays a role in the observed earnings premium. I also explore the causes of the lack of wage growth associated with AI skills via looking at the scarcity of AI skills, the connection between the earnings premium and the level of demand for AI skills, as well as wage effects of workers who could not have utilized the skill in previous jobs, i.e. recent graduates.

The baseline result in Section 4 showed that while there is an earnings premium, it turns non-significant when adding firm FE. The interpretation of this is that compared to peers within the same occupations, AI hires do experience an earnings premium, whereas compared to peers within the same occupations within the same firms, they do not. From this follows that the earnings premium should primarily stem from firms hiring AI workers paying all their employees higher wages.

The second result that warrants further investigation is the lack of differential wage growth when entering jobs requiring AI skills. This result is surprising, as AI skills are reported to be scarce and highly demanded. Being a barrier to the adoption of productivity increasing technologies, one would expect firms to use higher wages to poach workers with AI skills from their current employers. This lack of observed wage competition could have multiple explanations. First, it could be the case that despite reports of high scarcity, the supply of workers with AI skills actually meets the demand. This could, in turn, be a result of a number of factors. One possibility is that is that the productivity gains from using AI within a firm are not yet great enough to warrant the extra expenditure, and so firms may opt out of engaging in wage competition for these workers. This opens up the question of what will happen when AI technologies advance further. This, in turn, depends on how strong the connection between the AI earnings premium is to the demand for AI skills. With a strong connection, further advancements in AI technologies should then increase demand for AI skills. This should, all else being equal, increase the earnings premium.

The combination of a positive significant earnings premium with a lack of wage growth for individuals entering these positions also requires some consideration. The implication of a positive earnings premium in terms of level in the new job, combined with the lack of earnings growth associated with entering jobs requiring AI skills are illustrated in Figure 4. The lack of differential growth implies that these individuals already earned more than their peers in their previous employment. This opens up the question of whether individuals who are hired for these positions in most cases already utilized AI skills in their previous jobs, or whether these workers tend to be high ability individuals, who are moving between high paying jobs.

FIGURE 4
Earnings premium illustration

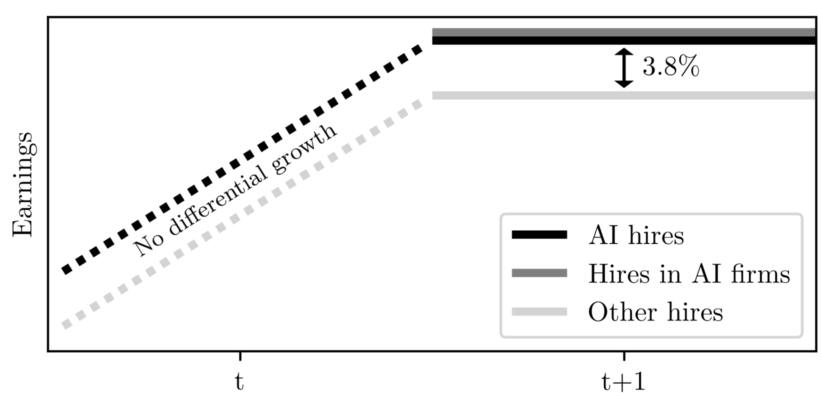

Notes: This figure illustrates a summary of the results, along with its implications. The earnings premium is taken from specification 4 of the baseline regression. The slope represents the median growth in earnings of job switchers between year $t-1$ and $t+1$. Median calculated as the mean growth of the 11 observations closest to the median.

5.1 Getting hired at a high wage firm

The observation that the earnings premium in the baseline regression disappears when controlling for firm fixed effects suggests that firms may not generally offer higher wages solely for employing an individual with AI skills. Rather, the AI skill premium might stem from an increased likelihood of securing employment in higher-paying firms. The nonsignificant firm FE regression results does not by itself prove this to be the case. First, splitting hires with AI skills into AI users and AI developers yielded a significant effect compared to peers within the same firms for AI developers. One could also imagine that the lack of significance when adding firm fixed effects simply stems from lack of variation in the data. This warrants further investigating to be able to make a firmer statement of whether the earnings premium can be explained by hires having sorted into high paying firms.

To explore this matching aspect of the earnings premium, I introduce a proxy for firm wage strategy in the form of AKM estimated firm fixed effects. In a regression using all employees (not just the new hires), I add person and firm-period fixed effects[20]. This approach is designed to isolate the impact of firm-level wage strategies within the firm-period fixed effects, with person fixed effects already controlling for occupation and traits influencing earnings. The percentiles of the resulting estimated firm-period fixed effects, $P(\hat{\psi}_{ft})$, serves as an outcome variable, with the objective to ascertain whether possessing AI skills is related to being hired in a firm with a certain wage strategy.

As the outcome variable is at the firm level, the regression cannot include firm fixed effects. Consequently, the specification including these (column (5) of the baseline regression) is excluded here.

Table 7 presents the outcomes of the firm wage outcome regression in columns (1) to (4). The first column suggests that AI skills are linked with a substantial 7 percentiles higher firm wage. However, this effect diminishes to 0.6 percentiles when accounting for occupation,

[20] Full details of the AKM regression outlined in Section F in the Appendix

TABLE 7

AKM firm wage percentiles

	Firm wage level				*Δ firm wage level*			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AI skill	6.677***	0.569	0.418*	0.421*	-0.419	1.115**	0.201	0.150
	(0.463)	(0.380)	(0.252)	(0.242)	(0.614)	(0.562)	(0.475)	(0.472)
Obs.	81,122	81,122	81,122	81,122	51,024	51,024	51,024	51,024
Adj. R^2	0.242	0.644	0.773	0.780	0.024	0.090	0.266	0.269
Control variables:								
Basic individual variables	✓	✓	✓	✓	✓	✓	✓	✓
Firm size				✓				✓
Fixed effects:								
Occupation		✓	✓	✓		✓	✓	✓
Workplace industry			✓	✓			✓	✓

Notes: This table displays estimates of regressions with two different outcome variables. In the left panel, the outcome variable is the estimated firm wage percentiles ($P(\hat{\psi}_{ft})$), and in the right panel the change in estimated firm wage percentiles ($P(\hat{\psi}_{ft}) - P(\hat{\psi}_{f,t-1})$)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. Standard errors in parantheses are clustered at occupation*firm*year level. *p<0.1; **p<0.05; ***p<0.01.

while only being significant at the 10% level in subsequent columns. I interpret these results as indicative of AI-skilled workers being associated with higher-paying firms, although there is not enough evidence to confirm that this is related to the AI skills, rather than the occupations these workers typically work in being more common in high paying firms.

However, the level of firm wage alone does not provide clarity on whether an individual with AI skills transitions to a higher-paying firm or merely shifts between two similarly high-paying firms. To delve into whether individuals equipped with AI skills are more likely to migrate to higher-paying firms, a separate regression makes the dependent variable the difference in the firm-period fixed effects percentiles between an individual's previous and current employer[21]. The results using this outcome in columns (5) to (8) are more varied, with only the regression adding occupation FE being statistically significant.[22] However, the

[21] $P(\hat{\psi}_{ft}) - P(\hat{\psi}_{f',t-1})$

[22] If one is to trust the one significant positive coefficient, which goes away when adding industry fixed effects, the interpretation could be that AI skills are in high demand in certain high paying industries, and workers with AI skills are now moving into these sectors. A common move for workers being hired into positions requiring AI skills is, in fact, to move from technology sectors into higher paying finance sectors, see common industry movements in Table A4 in the Appendix. Notably, 82% of individuals possessing AI skills hired in the finance industry had previously been employed in roles outside of the sector.

lack of a consistent pattern in these results means there is no clear takeaway.

5.2 Scarcity of AI skills

While the primary focus of this study is not centered on the demand and scarcity of AI skills in Sweden, this topic is intertwined to skill returns, as the valuation of AI skills should be related to its scarceness. If AI skills are scarce, firms may find themselves compelled to offer a wage premium to attract such talent. The existence of a significant earnings premium would imply a pronounced excess demand for AI workers, while also suggesting that firms perceive AI as a catalyst for substantial productivity enhancements, justifying the higher cost. Conversely, the absence of a discernible earnings premium could be indicative of two possible scenarios. Firstly, it may suggest that AI skills have become sufficiently widespread, removing the need for firms to engage in competitive behaviour to obtain workers with these skills. Alternatively, it might imply that the perceived productivity gains attributable to AI do not merit additional wage expenditures. In this section I investigate what can be gathered about the state of demand and scarcity of AI skills in Sweden during the sample period.

Starting in 2016, the demand for AI skills has experienced a sharp increase in Sweden, as illustrated in Figure A1 in the Appendix. This surge is mirrored in the United States, as found by Alekseeva et al. (2021). However, the heightened demand alone does not necessarily indicate an excess demand for AI skills.

The scarcity of AI skills is, however, found to be explicitly stated by Swedish firms in a survey by SCB (2023a), where a lack of qualified employees emerges as one of the primary obstacles to the adoption of AI. This barrier holds true for both firms currently using AI and those encountering difficulties in its adoption. Further supporting this notion, the study reveals that the least frequently cited barrier to AI adoption is the absence of perceived need within the firm.

FIGURE 5

Skill effects on successful hiring

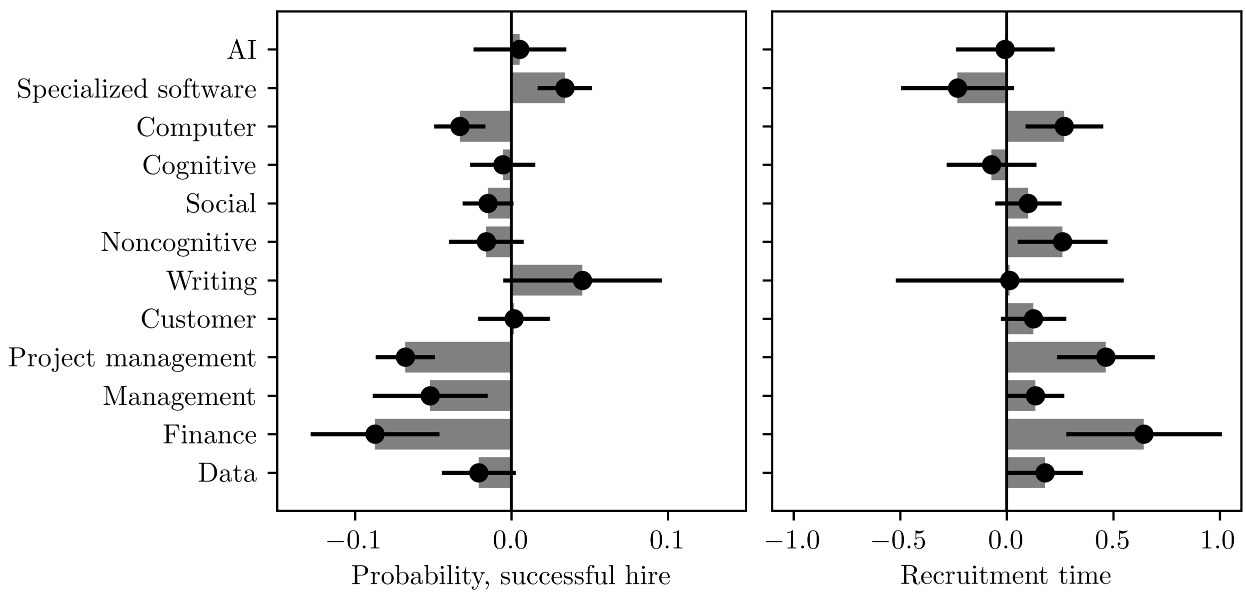

Notes: This figure displays the coefficients of a regression at the vacancy level, where the outcome variable is a dummy variable indicating whether the vacancy led to a hire (Obs. 1,823,090) as well as, in cases of success, how many months passed until the hire (Obs. 526,555). The explanatory variables are AI skills, data skills, and the Deming and Kahn (2018) skill categories. The regression controls for year and firm fixed effects, specification 3 and 7 of Table A5 in the Appendix.

I investigate the scarcity of AI skills myself by examining how the success of a recruiting process is related to the skilld required in the vacancy. This analysis aims to bridge the gap between the observed demand for AI skills and their actual scarcity. The underlying hypothesis is that difficulties in hiring could signify a genuine shortage of these skills. To investigate this, I run regressions at the vacancy level with the success rate of job vacancies and recruitment duration as dependent variables, while considering different skills, along with year and firm fixed effects, as explanatory variables. I define success as whether there was a hire matching the posted vacancy in the 5 months following the hire, and the recruitment time is the number of months elapsed between the vacancy posting and the start date of the first matching hire, limited to the sample of successful hires.

The findings of this analysis are presented in Figure 5. The regressions seem to uncover relations between skills and hiring difficulties well, given negative correlation between the likelihood of successful recruitment and the length of the recruitment process of most skills.

However, contrary to expectations, including AI skills in a job vacancy does not appear to be significantly related to either the probability of successful hiring or the recruitment duration. Additional specifications reported in Table A5 in the Appendix also do not reveal any connection between recruitment success and AI skills. The same pattern persists when excluding the other skill categories to prevent multicollinearity.[23]

While this outcome aligns well with the lack of wage growth found in Section 4.3, it seems at odds with the findings of SCB (2023a), which highlighted the scarcity of AI skills, especially as even firms that managed to adopt AI report having had difficulties in their recruitment processes. It is also unlikely that the rapid increase in demand for AI skills are being met by the supply of workers with these skills. I revisit this question in Section 5.4 to investigate whether the observed result may be explained by a lack of perceived benefit, making firms opt out of the recruitment process altogether.

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

[23]see Table A6 in the Appendix.

TABLE 8
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.

The results, shown in Table 8, 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[24].

5.4 *Increasing demand for AI skills*

As outlined in Section 5, 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, espe-

[24]See graph of experience and the earnings residuals of the baseline regression in Figure A6 in the Appendix

cially 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.

TABLE 9
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%.

6 Conclusion

Recent reports highlight a shortage of skilled workers as a key barrier to the adoption of AI technologies. I investigate whether this reported scarcity manifests in remunerations associated with being hired into a job requiring AI skills, as well as the drivers of this earnings premium. To investigate this, I use a method of attributing skill requirements from job vacancies to the individuals hired to fill the positions.

The results reveal that individuals employed in jobs requiring AI skills have an average earnings premium of 5%, 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.

I find some evidence indicating a potential discrepancy in the compensation of workers with AI skills contingent on whether they function as dedicated AI developers or merely apply pre-developed models. The former group seeing a higher compensation.

However, when comparing individuals with AI skills against their peers within the same firm, no additional compensation is discerned for possessing such skills. This suggests that another possible driver of the earnings premium is high wage firms employing individuals with AI skills. Despite this, a look into the wage strategies of AI hiring firms does not reveal any strong evidence in favour of this conclusion.

Despite the higher earnings experienced by workers with AI skills, it is revealed that these individuals do not experience any additional earnings increases upon entering positions requiring AI skills. Instead, the higher earnings seems to have carried over from past employment. This possibly indicates that most individuals with AI skills already utilized these skills in their previous employment, and that they now move between high-wage firms. This result also suggests that there is no significant wage competition among firms for workers with these skills, as firms do not seem to be willing to pay a premium in order to attract workers from their current employers. Consequently, this lack of wage competition could suggest that the high cost of workers with AI skills is, over the sample period, not a barrier to AI adoption in the Swedish economy.

However, the reasons behind this lack of wage competition remain unclear. Potential explanations range from sufficient supply of AI-skilled workers to firms opting for alternative methods of AI adoption such as outsourcing tasks or reskilling existing employees. It is also plausible that some firms simply do not perceive the benefits of AI adoption as outweighing the costs of hiring specialized AI workers, thereby opting out from starting the recruitment process altogether.

This wider question of AI adoption would benefit from additional research outside the scope of my paper. On the supply side, additional research into how and when individuals acquire AI skills could provide further context about the inflow and supply of workers with AI skills. On the demand side, exploring the motivations behind firms' decisions whether to adopt AI or not could help determine whether the absence of competition for workers with AI skills stems from alternative adoption strategies or from a lack of perceived benefits. This leaves the door open for some interesting future research topics. Moreover, considering the recent advancements in generative AI, and the evidence suggesting that the demand for AI skills is a driver of the AI earnings premium, it would be intriguing to revisit this question in the future, as more data from this period becomes available.

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., & Hodson, J. (2024). Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics*, *151*, 103745. doi: https://doi.org/10.1016/j.jfineco.2023.103745

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

Bartel, A. P., & Borjas, G. J. (1981, March). Wage Growth and Job Turnover: An Empirical Analysis. In *Studies in Labor Markets* (p. 65-90). National Bureau of Economic Research, Inc.

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

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

Cameron, A., & Miller, D. (2015, 03). A practitioner's guide to cluster-robust inference. *Journal of Human Resources*, *50*, 317-372. doi: 10.3368/jhr.50.2.317

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

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

Foster, A. D., & Rosenzweig, M. R. (2010). Microeconomics of technology adoption [Journal Article]. *Annual Review of Economics*, *2*(Volume 2, 2010), 395-424. doi: https://

doi.org/10.1146/annurev.economics.102308.124433

Frederiksen, A., Halliday, T., & Koch, A. K. (2016). Within- and cross-firm mobility and earnings growth. *ILR Review*, *69*(2), 320–353.

Hersch, J. (1998). Compensating differentials for gender-specific job injury risks. *The American Economic Review*, *88*(3), 598–607.

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

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

Lane, M., & Williams, M. (2023). *Defining and classifying ai in the workplace* (Working Paper No. 290). OECD Publishing. doi: https://doi.org/https://doi.org/10.1787/5f65ff7e-en

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

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/

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

Descriptive statistics, hire-vacancy sample

	Obs.	Mean	Median	Std.	Freq., = 1
Full matched sample:					
Mean AI skill	88,440	0.027	0	0.133	
AI skill = 1	88,440	0.014	0	0.117	1,226
AI skill > 0	88,440	0.077	0	0.266	6,778
AI skill >= 0.5	88,440	0.023	0	0.150	2,045
Number of matched vacancies	88,440	7.527	2	23.657	
Earnings, $t+1$	88,440	422.688	384.300	233.332	
Log(Earnings, $t+1$)	88,440	5.914	5.951	0.555	
Gender (1 = Female)	88,440	0.435	0	0.496	
Age	88,440	34.722	32	10.731	
Years of experience	88,440	10.877	9	8.435	
Number of employees	88,440	3,187.284	892	5,067.495	
Year	88,440	2,017.645	2,018	1.835	
Perfectly matched sample:					
AI skill = 1	40,156	0.026	0	0.160	1,061
Earnings, $t+1$	40,156	410.723	< mean	211.761	
Gender (1 = Female)	40,156	0.445	0	0.497	
Age	40,156	35.121	33	10.753	
Years of experience	40,156	11.533	10	8.467	
Number of employees	40,156	2,571.999	593	4,463.691	
Year	40,156	2,017.348	2,017	1.891	
Growth sample:					
AI skill >= 0.5	28,968	0.022	0	0.148	649
Earnings, $t+1$	28,968	454.793	< mean	224.322	
Earnings, $t-1$	28,968	370.746	357.300	240.747	
Earnings growth	28,925	0.302	< mean	0.680	
Gender (1 = Female)	28,968	0.434	0	0.496	
Age	28,968	39.180	38	10.351	
Years of experience	28,968	15.863	17	7.354	
Number of employees	28,968	2,921.103	< mean	4,884.430	
Year	28,968	2,017.160	2,017	1.691	

Notes: This table shows descriptive statistics of key variables in the matched hire-vacancy sample. The full sample is restricted in accordance with this baseline regression, and the growth sample with the earnings growth regression. Earnings is annual earnings in thousands of SEK. An observation equals to one hired individual. t is the year of hire. Medians consisting of micro-data replaced by size of median in comparison to mean.

Table A2

AI earnings growth premium, STEM interaction

	STEM education			*STEM occupation*		
	(1)	(2)	(3)	(4)	(5)	(6)
AI skill	-0.004	-0.023	-0.030	-0.001	-0.015	-0.025
	(0.035)	(0.034)	(0.038)	(0.025)	(0.024)	(0.026)
STEM education	-0.020	-0.023	-0.028*			
	(0.015)	(0.015)	(0.015)			
STEM education * AI skill	0.012	0.023	0.025			
	(0.040)	(0.040)	(0.042)			
STEM occupation * AI skill				0.032	0.036	0.060
				(0.052)	(0.052)	(0.056)
Obs.	20,978	20,978	20,978	20,978	20,978	20,978
Adj. R^2	0.135	0.140	0.171	0.135	0.140	0.171
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 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.

Table A3

Frequency table of AI hires in subgroups

	Level			*Growth*	
	STEM education	STEM occupation	Recent graduate	STEM education	STEM occupation
0	884	1785	1609	290	569
1	1161	260	436	359	80

Notes: This table displays the number of hires with AI skills with and without STEM education and occupations. Samples limited to those used in the level and growth regressions

Table A4

Industry movement of AI hires

Former industry	New industry	Share of hires	Earnings growth
Computer consultancy activities	Other monetary intermediation	2.9	7.7
Computer consultancy activities	Computer programming activities	2.6	2.8
Wired telecommunications activities	Wireless telecommunications activities	2.5	7.1
Computer programming activities	Other monetary intermediation	1.8	4.9
Computer programming activities	Computer consultancy activities	1.6	-2.8
Business and other management consultancy activ...	Computer programming activities	1.3	25.3
Manufacture of tubes, pipes, hollow profiles an...	Manufacture of tools	1.2	-7.5
Data processing, hosting and related activities	Other monetary intermediation	1.0	6.6
Engineering activities and related technical co...	Computer programming activities	0.8	23.3
Other software publishing	Computer programming activities	0.8	3.7
Other occupation changes		70.8	2.0
Total average growth		87.3	1.8

Notes: This table displays the ten most common combinations of industry changes among AI hires. The table excludes workers who did not change industry. Mean industry earnings growth is the differences in mean earnings of the two industries, among all employees in the period 2014-2020

Table A5

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 A6

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

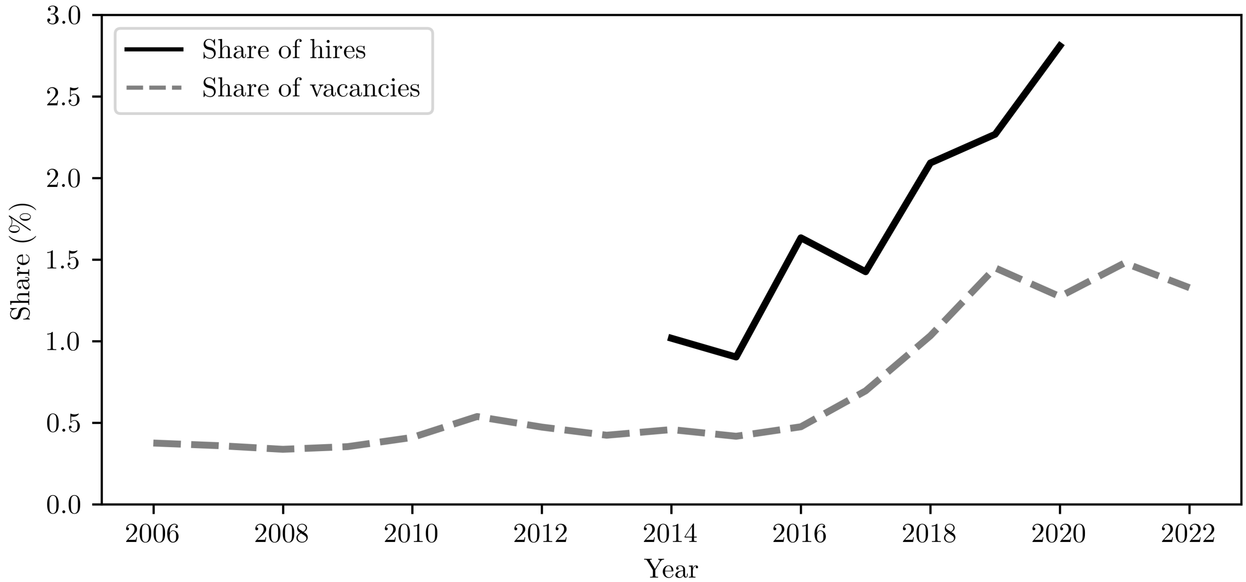

Notes: This figure displays hired AI workers as a share of total matched hires, and AI vacancies as a share of total number of vacancies from workplaces with at least one matched hire.

Figure A2

Sample comparison

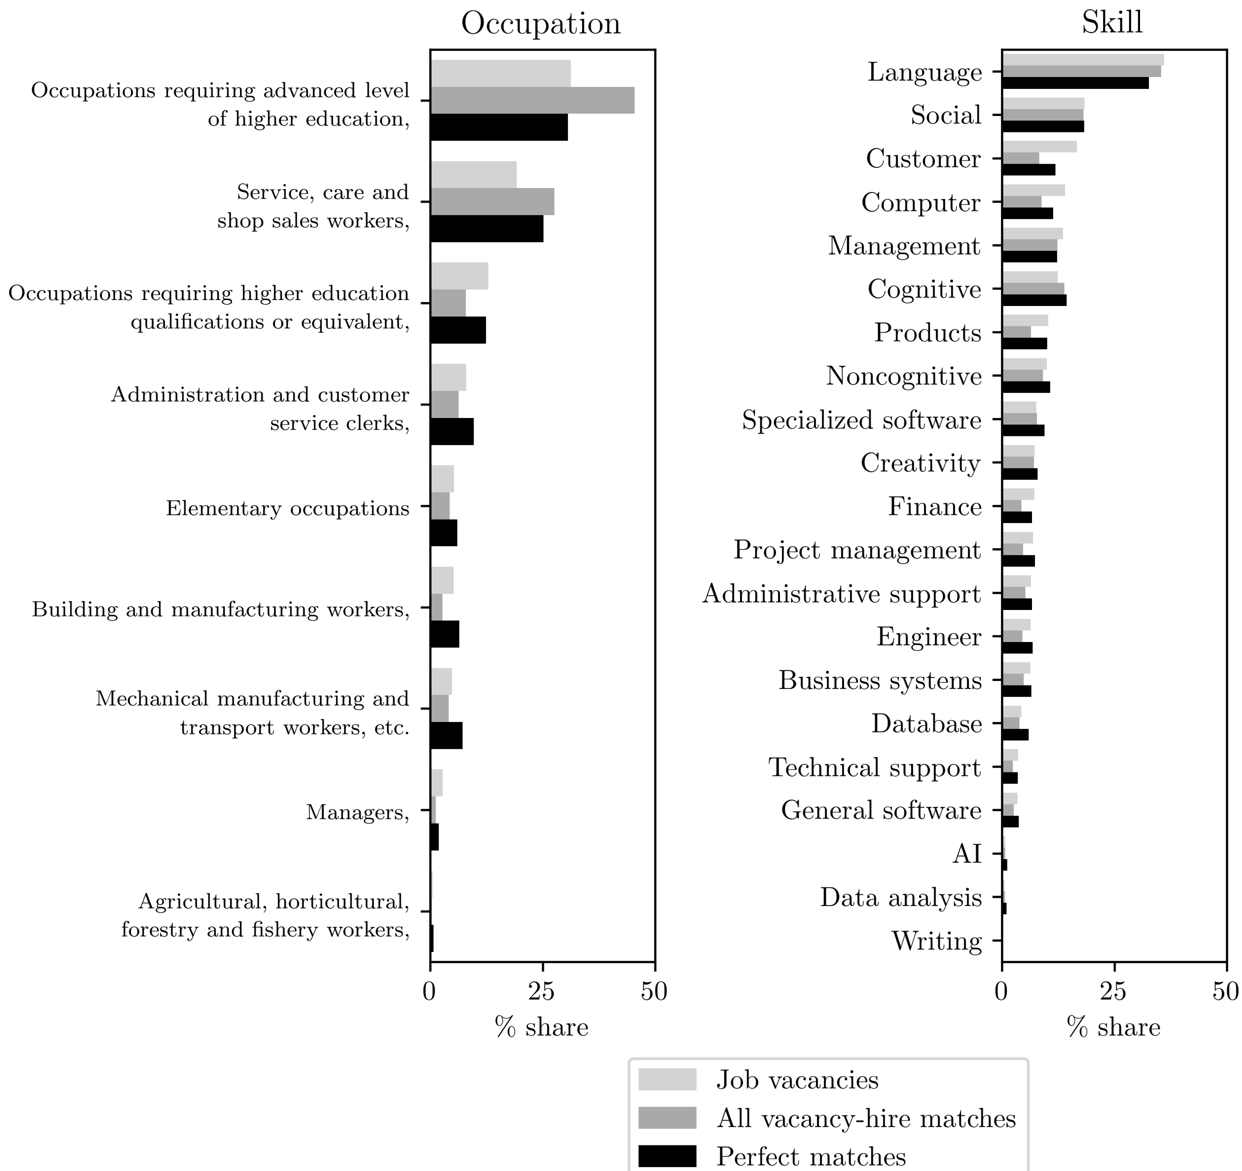

Notes: This figure is a comparison of differences between the posted vacancies, the vacancies that are in some way matched to a hire, and the sample of perfect vacancy-hire matches. The panel on the left displays occupations, at first level ISCO, and the panel on the right displays skills

Figure A3

Earnings growth distributions

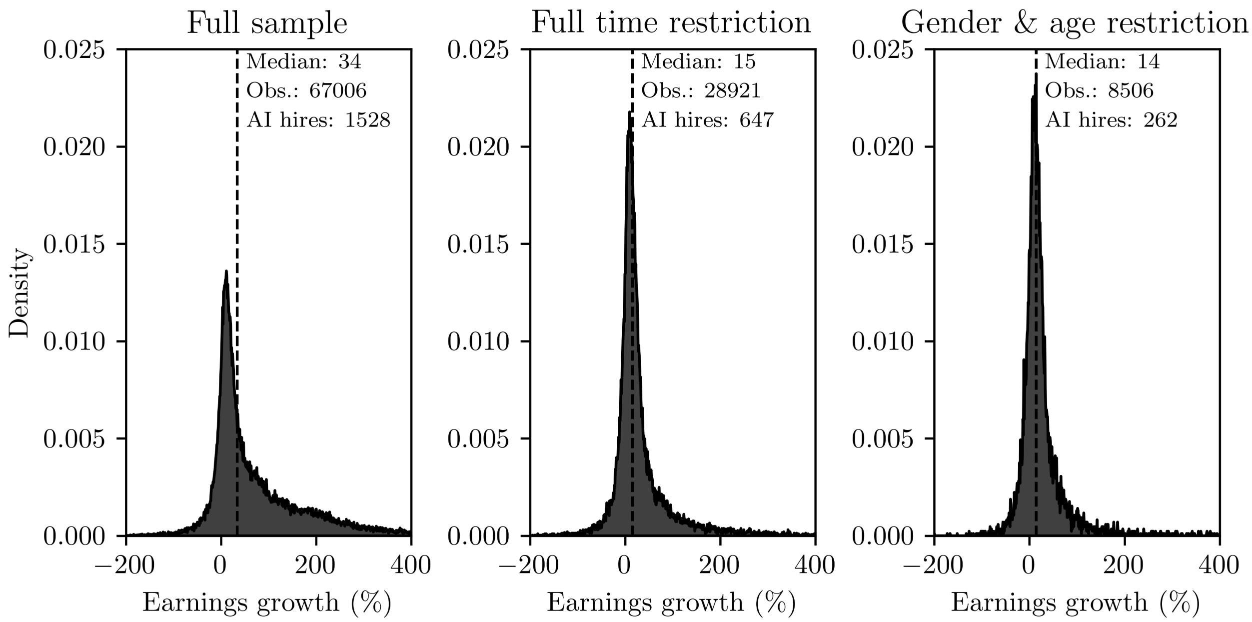

Notes: This figure displays the distribution of earninings growth at different sample restrictions. The full sample applies the basic sample restrictions outlined in section 2.2. The full time restriction additionally exluces workers who were not employed for the full calendar year t-1, workers with 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

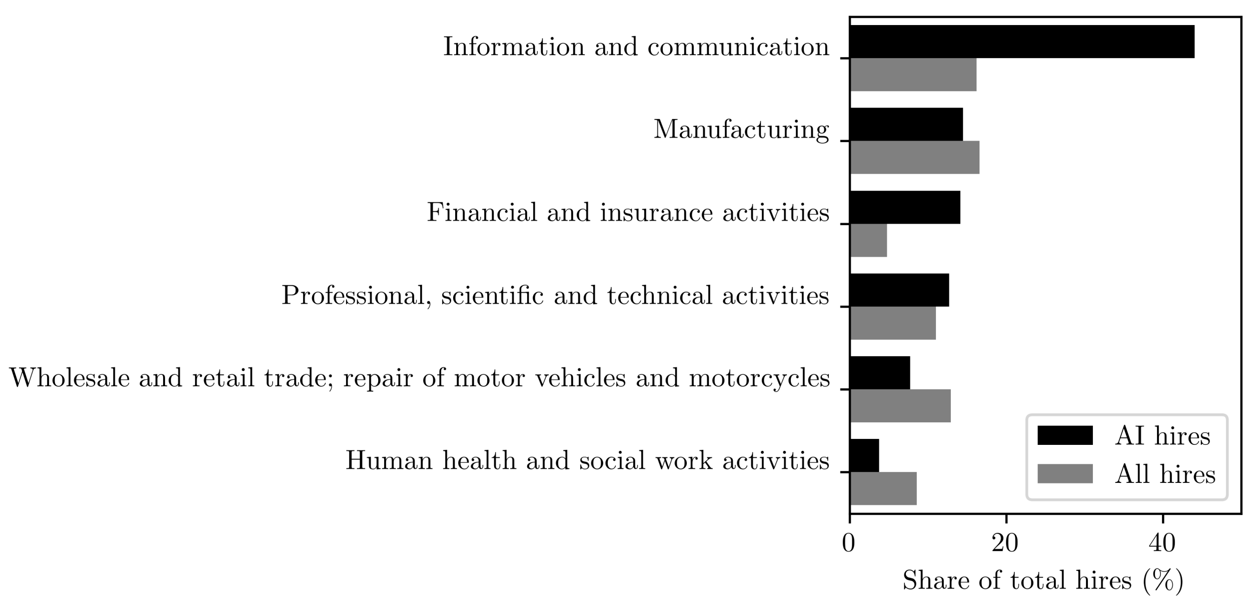

Notes: This figure displays AI hires by industry, as a share of all AI hires, compared to all hires. Industries at the first level NACE. Industries with less than one percent of all AI workers excluded.

Figure A5

Firm number of hires

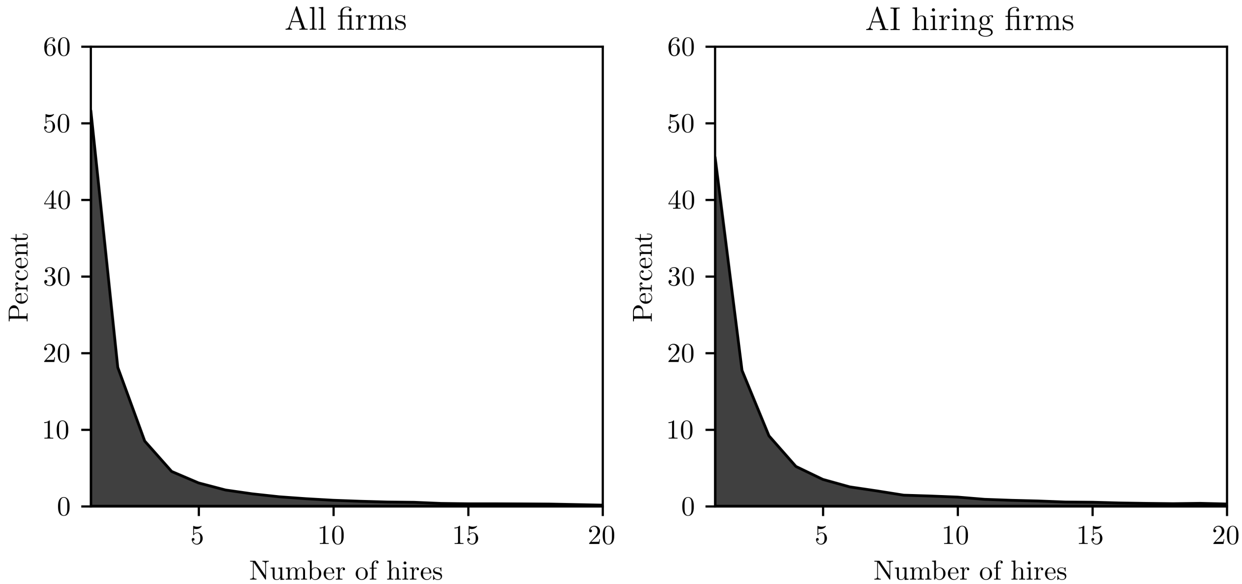

Notes: This figure displays histograms of the number of hires at the firm level, split into all firms, and firms with at least one AI hire. Observations with 5 or less firms are redacted.

Figure A6

Earnings and experience

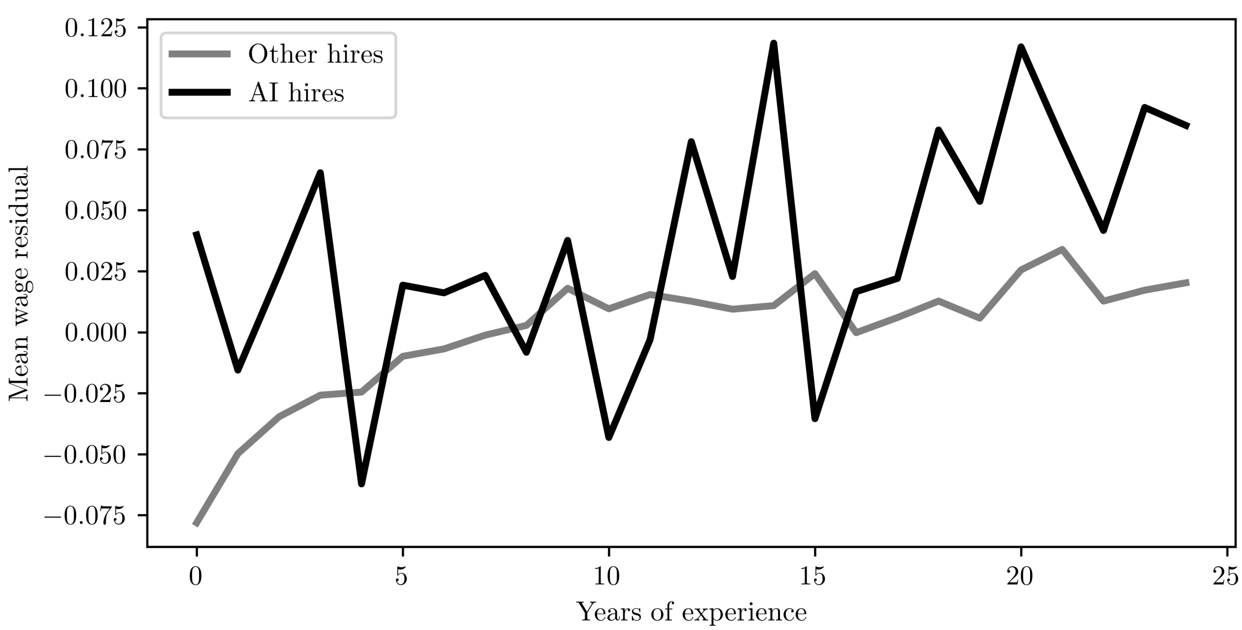

Notes: This figure displays the relationship between the natural logarithm of annual earnings, and experience, split by hires with and without AI skills. The y axis is the mean residuals of a regression based on specification 4 of the main regression, excluding experience, and experience-squared

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 A7.

Table A7
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 estimation 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. 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.

Figure A7
Hiring date approximation

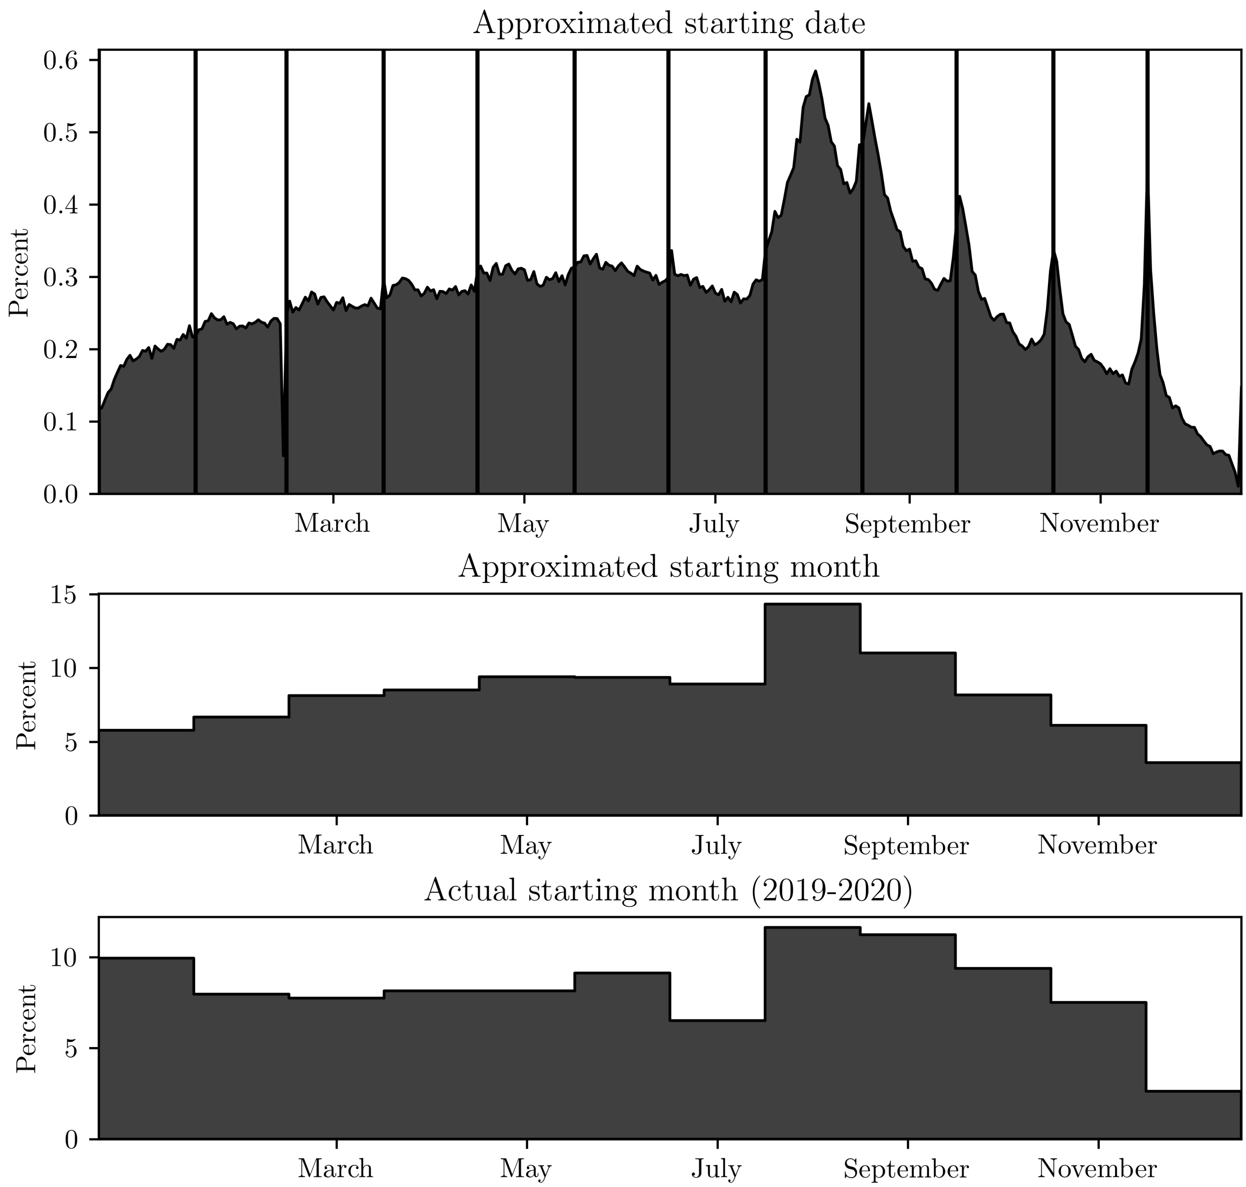

Notes: This figure displays the approximated hiring dates. The top graph is the date approximated by the share of earnings. The second panel turns the approximated date into hiring month, which is used for the final matching to job vacancies. The last panel displays the monthly distribution as recorded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

The graph in Figure A7 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.

E Robustness checks

E.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 A8 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 4 with varying cutoff values. The coefficients, displayed in Figure A8, reveal that the coefficient remains largely consistent and significant, irrespective of the chosen cutoff value.

Table A8

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.

E.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.

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

Figure A8

AI earnings regression, different cutoffs

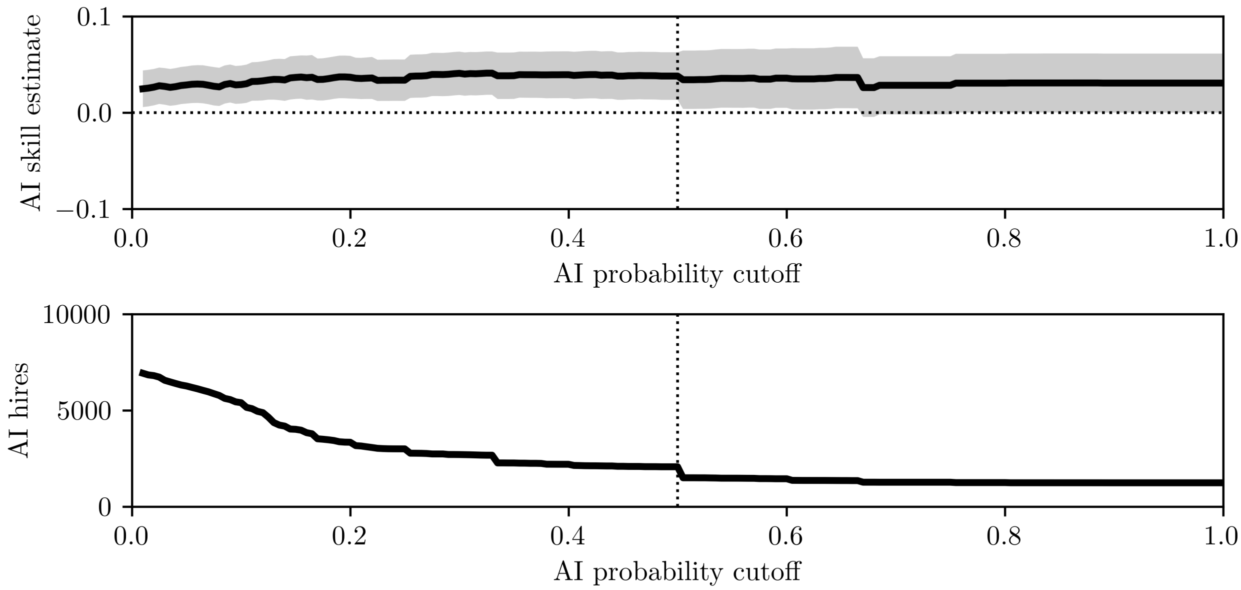

Notes: This figure displays the estimated effect of AI skill on earnings using different AI probabilities (I.e. mean value of skill of the matched vacancies of the hire.). Based on specification 4 from the baseline regression. Vertical line indicates the AI probility chosen for the main results.

Figure A9

Skewness of earnings residuals

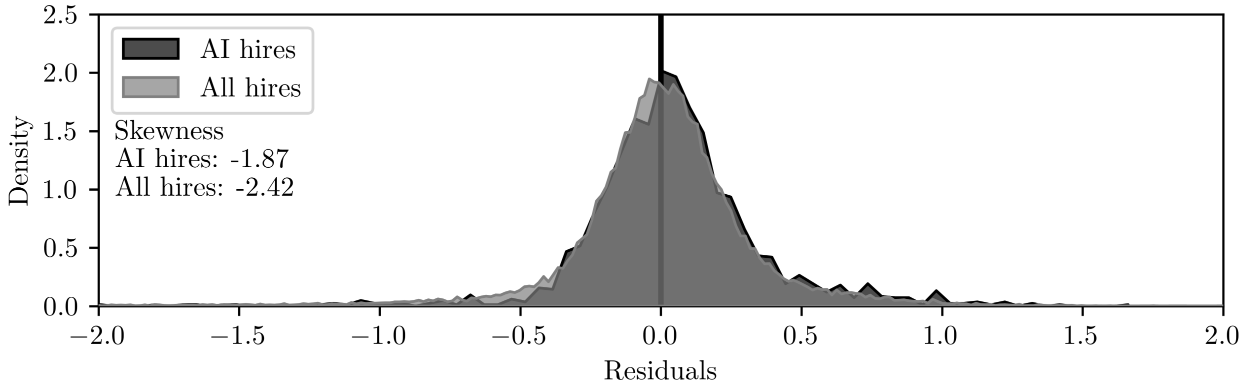

Notes: This figure displays the distributions of the residuals of a regression based on specification 2 from the main earnings premium result, with the sample limited to occupations with at least five AI hires. The dark histogram represents the residuals of AI hires, and the light gray histogram displays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

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 4. 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 A9. 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.

Table A9

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.

TABLE A10

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.

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 A9. 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 A10. 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.

E.3 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 A10, 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 A10
Skill correlation

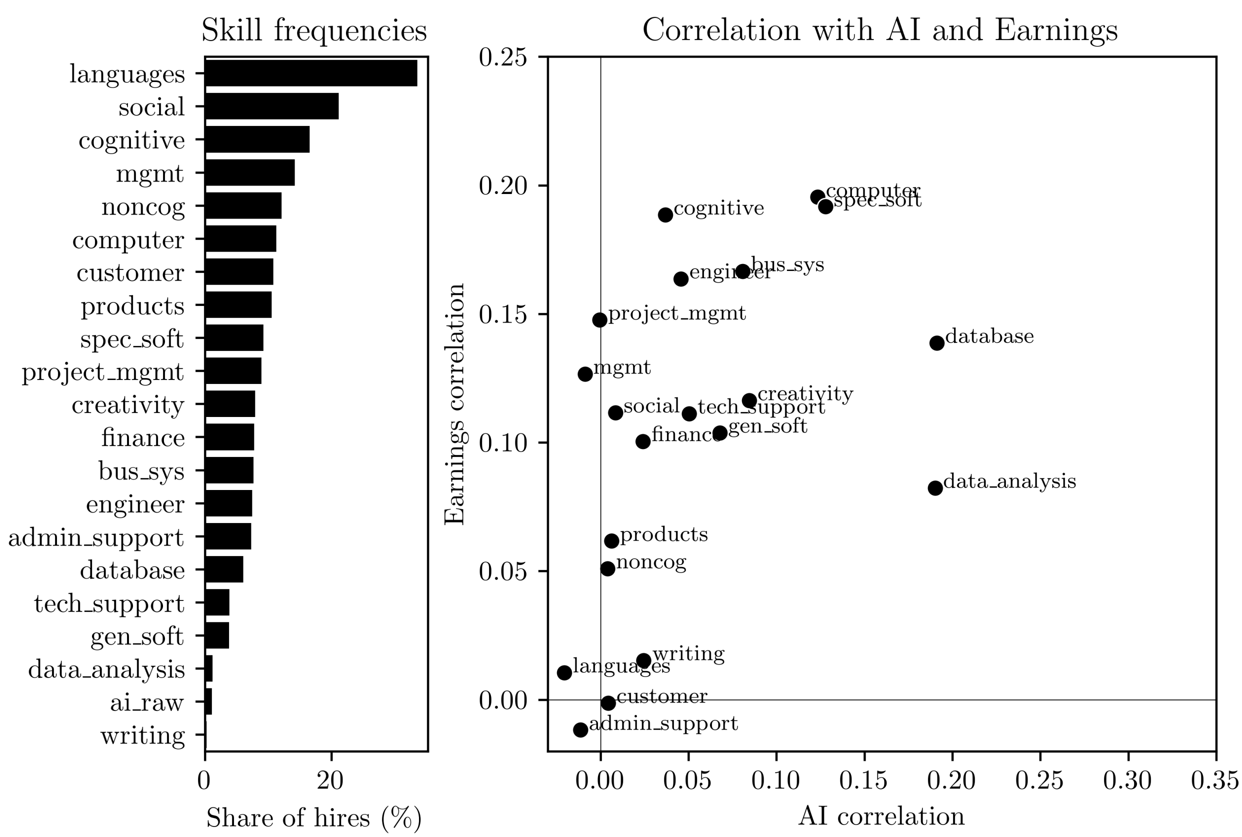

Notes: The panel on the left displays skill frequences in the perfectly matched sample of hires. The panel on the right displays spearman correlation between the Deming and Noray (2020) skill categories and AI skills, as well as with the natural logarithm of earnings.

Table A11 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 A11
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
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				
					(0.025)	(0.026)
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[25]. 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.4 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 much is the differing methodology. In order to investigate this, Table A12 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

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

to the findings of Alekseeva et al. (2021) when only controlling for firm fixed effects, but similarly to the basic earnings regression in Table 4, the combination of firm and occupation fixed effects removes the effect of AI skills.

Table A12
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.

E.5 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 A13. The results do not differ significantly from the baseline results.

Table A13

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.

E.6 Unreasonable earnings growth rates

As mentioned in section 3, there are some extreme growth rates left in the sample, despite the steps taken to mitigate this issue. A final safeguard against the remaining unreasonable growth rates driving the results, I impose a earnings growth cutoff. However, to avoid choosing an arbitrary cutoff value, I instead run a robustness check involving an iterative process of increasing sample restrictions on the top and bottom percentiles of earnings growth.

The graph in Figure A11, illustrates the coefficient of Specification 2 from the earnings growth regression. There is a consistent null effect irrespective of the proportion of excluded extreme values. There was some slight positive coefficients in the main growth regression, but this seems to be attributed to these outliers, as it converges towards zero upon excluding the top five percentiles.

Figure A11

AI earnings growth regression, different sample restrictions

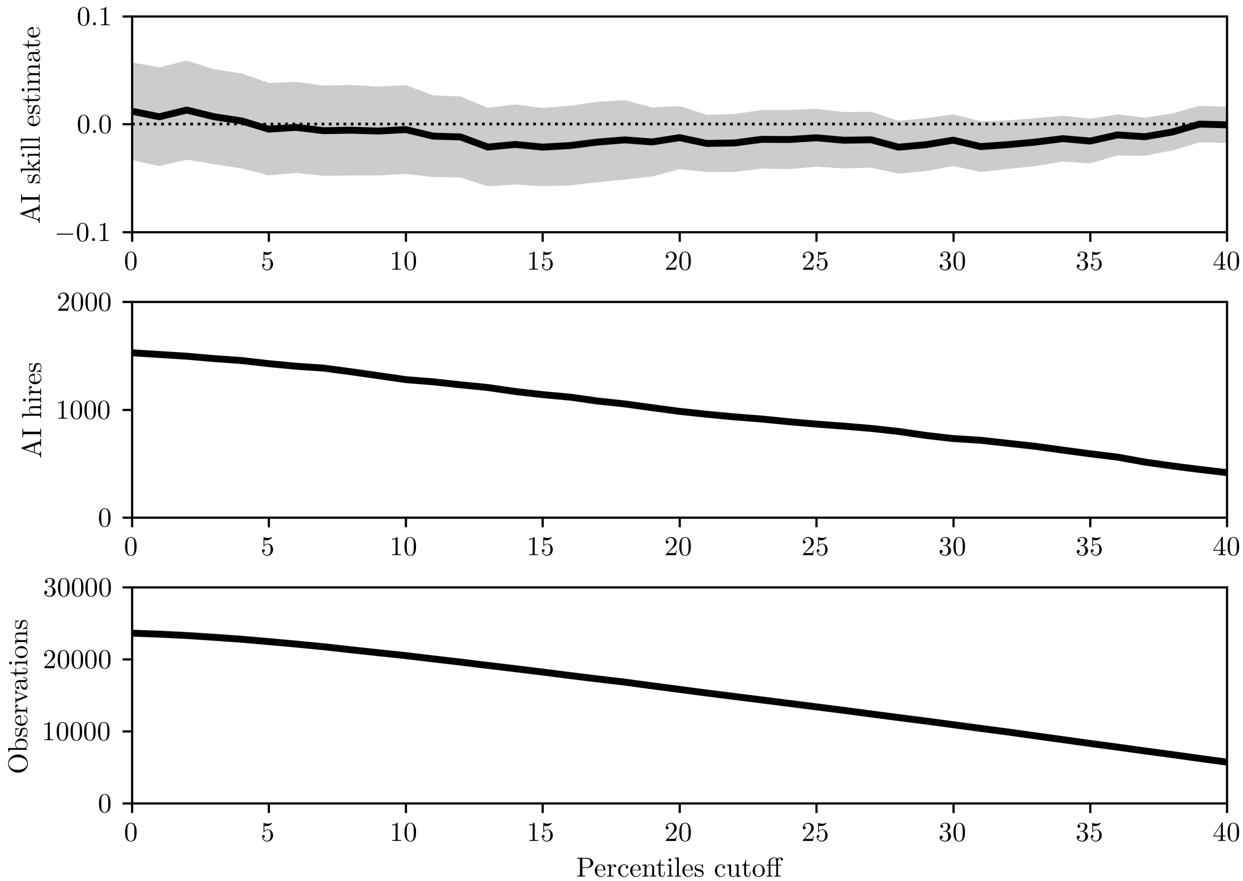

Notes: This figure displays the estimated effect of AI skill on earnings growth, using specification 2 from the earnings growth regression, at increasing sample restrictions. The top and bottom percentiles(Before other sample restrictions) are excluded from the sample according to the cutoff on the x-axis.

F AKM estimation

While the primary regressions in this study are limited to newly hired individuals, data on all employees of all firms spanning the years 1990 to 2020 encompassing all employees are available. Consequently, an AKM regression is conducted on this distinct sample, and the resulting estimates are used in a regression on the new-hire sample. The model utilized is as follows:

$$ln(earnings_{ift}) = \beta_0 + \psi_{ft} + \alpha_{it} + C_t + \beta_1 Unempl_{i,t-1} + \beta_2 Multi_{it} + \epsilon_{ift} \tag{2}$$

Here, ψ_{ft} denotes firm-period fixed effects, α_{it} signifies individual fixed effects, and C_t represents year fixed effects. As the data pertains to annual earnings, partial years of unemployment are not labelled, and would appear as lower earnings. To mitigate this, the inclusion of $Unempl_{i,t-1}$ is employed as a binary variable indicating unemployment status in November of the preceding year, and $Multi_{it}$ serves as a binary indicator denoting employment at multiple firms within the given year.

The rationale behind incorporating Firm-period fixed effects stems from the inquiry into whether individuals are prone to gravitate towards firms offering higher remuneration relative to a benchmark firm, holding all else constant. Given the potential for shifts in a firm's wage policy over time, firm-period effects are employed, with each period spanning 7 years, the final one encapsulating the main regression period of 2014-2020.

In pursuit of understanding wage dynamics for newly employed individuals, the individuals in the matched hiring sample are excluded from the AKM regression during the year of hiring and all subsequent years.

TABLE A14
AKM estimates output

	(1)	(2)
Multiple jobs	-0.033***	-0.066***
	(0.000)	(0.000)
Unemployed last year	-1.160***	-0.734***
	(0.000)	(0.000)
Obs.	150,104,173	137,871,289
Adj. R^2	0.598	0.721
% share of final sample excluded:		
Person	9.2	
Firm	0.0	

Notes: This table displays the estimates of the AKM regression. Specification 1 includes firm-period, person, and year fixed effects, the second specification replaces firm-period and person with person-firm fixed effects. % share of final sample excluded indicates the percentages of the final sample used in the main regressions that are excluded due to perfect colinearity.

The outcomes of this regression are displayed in Table A14. The observed R^2 value is low, at around 60%. Plausible explanations for this could encompass various factors influencing earnings, which would not be captured in a pure wage variable, such as transitions between jobs, part-time work, incentive-based compensation, among others. In order to check the performance when including the most detailed variables available, an alternate specification (specification 2) is introduced wherein person and firm fixed effects are substituted with person-firm fixed effects. This alteration results in an R^2 of 72%, which is still relatively low. Consequently, I advise caution when interpreting results based on the AKM findings in the main results.