

STOCKHOLM SCHOOL OF ECONOMICS  
DEPARTMENT OF ECONOMICS  
BE551 DEGREE PROJECT IN ECONOMICS  
SPRING 2026

# AI and Junior Hiring – Evidence from Swedish Job Postings

Alexander Ekdahl (26134) and Rasmus Samuelsson (25969)

**Abstract:** The rapid diffusion of Generative Artificial Intelligence (GenAI) has raised widespread concerns about its employment effects for workers at the earliest stages of their careers. In this thesis, we investigate whether Generative AI has disproportionately reduced the number of vacancies for junior workers, who typically perform the routine tasks that AI is best suited to replace. Using approximately 3.6 million job postings from Sweden's largest job boards spanning 2021 to 2026, we classify postings by seniority using a novel keyword-based classification of job title and implement a fixed-effect difference-in-differences design that exploits variation in occupational AI exposure before and after the release of ChatGPT in November 2022. We find evidence consistent with higher AI exposure occupations experiencing a disproportionate decline in junior share of seniority-classified postings following the release of ChatGPT, with the effects emerging gradually from 2024 onward.

**Keywords:** AI, Junior Hiring, Entry-Level Employment, Labor Market, ChatGPT, Difference-in-Differences, Sweden

**JEL:** J23, J24, J63, O33, C23

|                 |                     |
|-----------------|---------------------|
| Supervisor:     | Álvaro Janez Garcia |
| Date submitted: | 11 May 2026         |
| Date examined:  | 18 May 2026         |
| Discussant:     | Julia Cook          |
| Examiner:       | Johanna Wallenius   |

## **Acknowledgements**

First, we would like to extend our gratitude to our supervisor Álvaro Janez Garcia, for his great advice and insights throughout the thesis writing process. We would also like to thank our classmates, friends and family for the parts they have played in completing this work.

## Abbreviations

|       |                                                    |
|-------|----------------------------------------------------|
| AI    | Artificial Intelligence                            |
| DAIOE | Dynamic AI Occupational Exposure                   |
| DiD   | Difference-in-Difference                           |
| GenAI | Generative Artificial Intelligence                 |
| ISCO  | International Classification for Occupations       |
| LLM   | Large Language Models                              |
| OLS   | Ordinary Least Squares                             |
| SACO  | Swedish Confederation of Professional Associations |
| SBTC  | Skill-Biased technological Change                  |
| SOC   | Standard Occupational Classification               |
| SSYK  | Swedish Standard Occupational Classification       |

## Table of Contents

|          |                                             |           |
|----------|---------------------------------------------|-----------|
| <b>1</b> | <b>Introduction.....</b>                    | <b>1</b>  |
| <b>2</b> | <b>Literature Review and Theory .....</b>   | <b>3</b>  |
| 2.1      | Automation and the Labor Market .....       | 3         |
| 2.2      | Skill-Biased Technological Change.....      | 4         |
| <b>3</b> | <b>The Theoretical Framework.....</b>       | <b>6</b>  |
| 3.1      | Model Intuition .....                       | 7         |
| 3.2      | Comparative Static.....                     | 8         |
| 3.3      | Combined Effect .....                       | 9         |
| <b>4</b> | <b>Data .....</b>                           | <b>9</b>  |
| 4.1      | JobTech Links.....                          | 10        |
| 4.2      | AI Exposure Indices .....                   | 10        |
| 4.3      | Summary Statistics.....                     | 11        |
| <b>5</b> | <b>Empirical Strategy .....</b>             | <b>13</b> |
| 5.1      | From Model to Empirical Specification ..... | 13        |
| 5.2      | Dependent Variable.....                     | 13        |
| 5.3      | Seniority.....                              | 14        |
| 5.4      | Cutoff.....                                 | 15        |
| 5.5      | Observational Window.....                   | 16        |
| 5.6      | AI Exposure .....                           | 17        |
| 5.7      | Control Variables .....                     | 17        |
| 5.8      | Model Specification .....                   | 17        |
| 5.9      | Identifying Assumptions .....               | 19        |
| <b>6</b> | <b>Results .....</b>                        | <b>20</b> |
| 6.1      | Fixed Effect OLS – Within Occupation.....   | 20        |
| 6.2      | Fixed Effects OLS – Within Firm .....       | 21        |
| 6.3      | Alternative Specifications .....            | 24        |
| 6.4      | Robustness Checks .....                     | 27        |
| <b>7</b> | <b>Discussion.....</b>                      | <b>30</b> |
| 7.1      | Interpretation of Findings.....             | 30        |
| 7.2      | AI Index Heterogeneity .....                | 33        |
| 7.3      | Limitations .....                           | 33        |
| 7.4      | External validity .....                     | 36        |
| 7.5      | Comparison to Other Studies .....           | 36        |

|            |                                                                          |           |
|------------|--------------------------------------------------------------------------|-----------|
| <b>8</b>   | <b>Conclusion .....</b>                                                  | <b>37</b> |
| <b>9</b>   | <b>Bibliography .....</b>                                                | <b>39</b> |
| <b>10</b>  | <b>Appendix.....</b>                                                     | <b>42</b> |
| <b>A.1</b> | <b>AI disclosure statement .....</b>                                     | <b>42</b> |
| <b>A.2</b> | <b>Data – Job postings.....</b>                                          | <b>42</b> |
| <b>A.3</b> | <b>Empirical strategy – Seniority keywords .....</b>                     | <b>43</b> |
| <b>A.4</b> | <b>Theoretical Framework – Derivation of the comparative Static.....</b> | <b>44</b> |

# 1 Introduction

After its launch on November 30th 2022, ChatGPT, an LLM (Large Language Model) created by OpenAI, reached 1 million users in only 5 days, 100 million in 2 months (Hu, 2023). As of February 2026, the service is estimated to draw over 900 million weekly active users, a trajectory that has prompted a near trillion-dollar valuation of OpenAI. Early estimates of the technologies' transformative potential are seemingly in line with these valuations. Eloundou et al. (2024) estimate that 80 percent of the U.S. labor force could have 10 percent or more of their tasks affected, with 19 percent of workers seeing at least 50 percent of their tasks impacted. While technological change has historically generated long-run productivity gains, it has also produced significant short-run employment losses in the sectors and occupations most directly exposed, and generative AI's effect on junior workers is increasingly seen as no exception.

There are early indications that this transformation is already underway. The rapid diffusion of generative AI has coincided with an increasingly competitive labor market for workers at the earliest stages of their careers. Unemployment rates for recent college graduates in the U.S. have risen sharply, despite stable youth unemployment overall (Hosseini et al. 2026), and job postings for entry-level roles requiring a degree in the United Kingdom have fallen by almost two-thirds since 2022 (Murray, 2025). If entry-level roles are increasingly susceptible to substitution from AI, the returns to educational investment become less certain. Junior positions have traditionally served not only as productive roles, but as the foundation of long-run career development. Their displacement raises broader questions about how workers enter and progress through the labor market.

This paper makes two contributions to the field. First, we introduce job posting titles as a novel means of measuring seniority, classifying vacancies by the keywords firms use to describe the role rather than the age or tenure of workers ultimately hired. Second, we contribute to an emerging literature with mixed international findings on whether generative AI is reshaping entry-level hiring across different institutionalized settings with different rates of technological diffusion.

We exploit the release of ChatGPT on November 30<sup>th</sup> 2022 as an exogenous shock to the productivity of generative AI. ChatGPT represented a discrete shift in the accessibility and capability of large language models, providing a clean temporal discontinuity around which to identify labor market responses. Occupations differ substantially in the degree to which their task content overlaps with current AI capabilities, providing cross-sectional variation in the treatment intensity. Our identification strategy combines two sources of variation, comparing the evolution of junior vacancy postings across occupations with heterogeneous AI exposure, before and after the release of ChatGPT.

We draw on the near universe of Swedish job postings between 2021 and 2026. The JobTech Links dataset, a collaboration between Swedish Public Employment Service and several of Sweden's largest job posting platform captures approximately 3.6 million postings from roughly 250,000 firms, providing comprehensive coverage of job vacancy activity in the

Swedish labor market. We classify the seniority of each posting by flagging job titles for seniority-coded keywords, yielding a measure of junior share, the fraction of seniority-classified postings designated as junior, at the occupation and firm level. Occupational AI exposure is measured using two distinct indices, Eisfeldt et al. (2023) and Engberg et al. (2023), capturing the extent to which an occupation’s task content corresponds to the AI capabilities during the time of the natural experiment.

We employ two complementary difference-in-differences specifications. The first estimates a fixed effects OLS model at the occupation-half-year level, identifying the effect from differential changes in junior share between high- and low-AI-exposure occupations over time, conditional on occupation and time fixed effects. The second estimates a fixed effects OLS model at the firm-half-year level, identifying the effect from differential changes in junior share across occupations, within the same firm in the same period, absorbing firm-level demand shocks that could otherwise confound the occupation-level estimates. We complement both specifications with event study estimates that test for pre-existing differential trends and characterize the timing of the effect relative to the release of ChatGPT.

In our preferred firm-level specification, which identifies the effect from differential within-firm reallocation across occupations rather than economy-wide trends, occupations with higher AI exposure experience a larger decline in junior share following the ChatGPT release. The effect is gradual rather than immediate, emerging clearly from 2024 onward and deepening through the end of our sample window, consistent with AI adoption operating through a diffusion process rather than a discrete shock. Decomposing the result into its constituent margins, we find the decline in junior share is driven by a reduction in junior postings rather than an expansion of senior postings, supporting the substitution channel and arguing against a purely macroeconomic interpretation. Notably, we find no significant relationship between AI exposure and senior postings, indicating that the broad decline in high-AI occupations is concentrated further down the seniority distribution.

The main results are robust to a series of checks. Excluding the most prevalent occupations, personal care assistants and software engineers, individually leaves the sign and significance unchanged in the firm-half-year FE specification. This suggests that results are not driven by any trend in occupational groups. The direction and timing of the effect are unchanged under a standard calendar half-year split, with the loss of precision consistent with seasonal contamination of the reference period. The sign of the interaction estimate is generally consistent across both AI indices, reducing concern that findings are an artifact of any specific exposure measure. The results are most robust and precisely estimated when restricted to high-skill professional occupations, where the theoretical assumption and validity of our seniority classification are most plausibly satisfied.

The remainder of the paper is structured as follows. Section 2 reviews the relevant literature. Section 3 presents the underlying theoretical framework. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 presents the main findings and robustness checks. Section 7 discusses the results and section 8 concludes.

## 2 Literature Review and Theory

This paper primarily relates to two strands of economic literature. The first concerns the broad relationship between technological development and labor market outcomes. The second is skill-biased technological change, how new technologies discriminate labor market participants along the dimension of competence.

### 2.1 Automation and the Labor Market

The effect of automation technologies on labor demand has been a central concern of labor economics, with research spanning centuries of industrial transformation. Acemoglu and Restrepo (2019) establish a foundational framework for understanding this relationship, decomposing the effect of automation into three competing forces: the *displacement effect*, whereby machines substitute for labor in existing tasks; the *productivity effect*, whereby automation lowers costs and raises output, increasing demand for labor elsewhere; and the *reinstatement effect*, whereby new technologies create entirely new tasks in which labor holds a comparative advantage. There is no pre-determined direction of the net effect of these forces, instead, it depends on the characteristics of the automation technology in question. The authors find evidence that over the preceding two decades, displacement has outpaced reinstatement in the U.S. economy, attributing the almost complete stagnation in labor demand growth over the last two decades in the U.S. to technology-induced shifts in the labor market caused by automation, challenging the presumption that technological progress is inherently favorable to labor.

The tension between these forces is observable across historical episodes of technological change. Examining wages and employment for textile-industry workers in 19<sup>th</sup> century Britain, Acemoglu and Johnson (2024) find that the introduction of spinning machines caused employment of hand-loom weavers to decrease by roughly 80 percent and almost halve real wages. Considering a more contemporary technological shift, Acemoglu and Restrepo (2020) study the adoption of industrial robots across the U.S. and find evidence that one additional robot per thousand workers in the U.S. reduces the employment-to-population ratio by 0.2 percentage points, and exerts a dampening effect on wages of roughly 0.4 percent.

To what extent Artificial Intelligence will raise productivity, create new jobs, or substitute for human labor remains an open question, but early evidence suggests that the scale of substitutional exposure is substantial. Eloundou et al. (2023) claims that 15 percent of all tasks could be completed significantly faster and with same level of quality, or higher, when utilizing an LLM, with 19 percent of workers seeing at least 50 percent of their tasks impacted. Whether this exposure has translated into realized labor market effects is contested. Acemoglu et al. (2022) study the impact of AI on labor markets using job postings in the pre-GPT era, finding that firm-level AI adoption reduces vacancies in non-AI positions. However, they find no relationship between occupational AI exposure and wages or employment at the occupation or industry level. Engberg et al. (2025) perform a similar analysis on Swedish postings between 2014-2022. Contrary to findings in the US, they observed a positive effect of AI exposure on both non-AI hiring and total employment in Sweden. The authors suggest that this shows a tendency for AI to augment rather than replace



non-AI work in a Swedish context. This highlights the possible discrepancy between what has been observed in the U.S. (Acemoglu, 2022) and what could be observed in a Swedish context. Chandar (2025) uses U.S. current Population Survey data following the release of ChatGPT and finds no difference in employment and earnings between high-AI exposure occupations and low-AI occupations. Bick et al. (2026) explore the discrepancies in AI adoption between countries. Most notably they find that 43 percent of U.S. workers used AI for their jobs in 2026, compared to a range of 26 to 36 percent for Germany, UK, France, Italy, Netherlands and Sweden. Bick et al. (2026) reason that the adoption gap is caused by firm management practices and employer encouragement to use AI. From European firm survey data between 2021 and 2025, they show that the AI use of firms from the forementioned countries was stable between 2021 and 2023, before more than doubling after 2023, lagging their US counterparts who accelerated in their adoption earlier. Taken together, these studies suggest that while the theoretical potential for substitution is large, aggregate employment effects have so far been limited, partly the result of a relatively gradual diffusion process, only accelerating over a year after the release of ChatGPT.

## 2.2 Skill-Biased Technological Change

The effects of technological change on the labor market are not uniform across workers. A central question in the research on automation concerns the heterogeneity of these effects, specifically whether new technologies systematically favor certain types of workers over others. Skill-biased technological change is a particular framework within the literature on technological effects on the labor market focusing on how automation shifts labor demand away from low-skill workers and towards those with higher levels of education and competence.

In one of the landmark papers within the field of labor economics, Autor, Levy, and Murnane (2003) find that occupational task structure, more precisely the degree to which a task can be considered routine, meaningfully influences the effect of automation technologies. The authors find that the computerization of the late 20<sup>th</sup> century is associated with a decrease in labor input for routine tasks and increase in labor input for nonroutine tasks, an effect the authors argue is caused by computer technologies substituting low-skill workers performing routine tasks, while complementing high-skill workers performing nonroutine tasks, raising the relative demand for skilled workers. They attribute computerization driven task-shifts for 60 percent of the relative demand shift favoring college-educated workers between 1970 and 1998. Dillender et al. (2022) study computerization in a more applied setting. They find that when firms adopt new software at the job-title level, skill requirements in job postings increase. The authors also show evidence of skill-bias, with wages for college graduates in office administration occupations increasing with negative wage spillovers in the rest of the local economy.

Research on the skill-biased effect of AI points toward it potentially being different from the example of computerization. Autor (2024) argues that AI can extend expertise to low-skill and middle-skill workers, providing them with access to knowledge that historically would require years of training to acquire. Scholars such as Webb (2020) go as far as to argue that AI would constitute a reversal of the traditional predictions of the skill-biased technological

change framework, arguing that high-skill workers are most exposed to AI and that it is likely to decrease 90-10 wage inequality. Many empirical examinations of AI are line with the prediction that the benefits of AI accrue to low-skill workers. Brynjolfsson et al. (2025b), find that the adoption of LLMs by call centers significantly increased the productivity of less skilled and less experienced workers, while having no effect on the productivity of more experienced workers. Noy and Zhang (2023) perform a Randomized Control Trial and find that access to ChatGPT decreased inequality between workers, with lower-skilled workers benefiting the most in professional writing tasks. While set within a relatively narrow set of tasks, these findings portray AI as a tool that increases the relative productivity of low-skill workers. However, the relationship between AI-driven productivity gains and demand for junior workers is ambiguous. Even if AI were to raise the productivity of less experienced workers, this does not unambiguously imply that firms will hire more of them. A productivity gain reduces the cost of a task, but if AI can perform the task directly, the firm may simply substitute away from the worker performing it all together.

The discriminatory effects of technological change along the dimension of skill raises the question of what other worker characteristics determine differential consequences of technological development. While the SBTC (Skill-Biased technological Change) framework examines the differential effect of technological development along the dimension of education and formal skill, a nascent strand of literature suggests that experience may be an equally important axis of differentiation. Senior workers are not simply more skilled versions of junior workers, but they possess accumulated tacit knowledge, judgement and professional relationships that are difficult to codify, and therefore difficult to automate. Rather than asking whether new technologies favor high-skill over low-skill workers, this literature asks whether they systematically disadvantage workers earlier in their careers, whose tasks tend to be more bounded and codifiable, relative to senior workers who leverage the experience they have accumulated over their careers.

Brynjolfsson et al. (2025a) found that early career workers experienced a 16 percent relative employment decline following the availability of ChatGPT, particularly concentrated in professions where AI enables automation rather than in those in which it augments labor, where the effect was the opposite. Westby and Modestino (2025) use software engineering job postings classified by experience requirements and find that generative AI resulted in a large and significant decrease in the ratio between junior and senior software developers, with junior postings falling by approximately 16 percent. Klein Teeselink (2025) employs a DiD design, comparing junior employment to senior employment within firms and occupations in the UK and find that AI exposed firms reduce employment, with a particularly pronounced effect in junior positions. In an attempt at establishing a more causal effect, Hosseini et al. (2026) goes beyond exposure measures, using job postings for AI integrators as a proxy for firm-level AI adoption and a detailed worker dataset with hierarchical seniority, and finds that junior employment in AI adopting firms decreases by about 10 percent relative to non-adopters 18 months after the release of ChatGPT. Humlum and Vestergaard (2026) employ surveys on firm encouragement of AI chatbot use and administrative labor record in Denmark to study the effect of AI adoption on a range of labor market outcomes. They replicate the decrease in employment for young workers in exposed occupations seen in Brynjolfsson et al

(2025a) but find no difference in trend between firms that encourage AI chatbot use and those that do not.

While previous studies on generative AI and the labor market and its effect on junior hiring have been performed in an Anglo-Saxon context, a limited number of studies have been performed with the Swedish labor market in mind. First, Engberg et al.’s *AI Unboxed and Jobs* (2024), who uses register data on firms and individuals between 2009 and 2023 and finds a positive association between AI exposure and labor demand for high-skilled white-collar workers, in line with the upskilling effect identified by Autor et al. above. However, the panel only covers one month after the release of ChatGPT, mostly focusing on pre-GPT AI use. Second, closest to our paper, and released as a contemporaneous working paper several months into our work, Engberg et al. (2026) find a sharp decrease in the employment of 22–25-year-olds in the occupations most exposed to AI. Worth highlighting is that our present paper was conceived and the empirical analysis largely completed independently and prior to the release of the Engberg et al. working paper in late March 2026.

### 3 The Theoretical Framework

The literature reviewed above motivates a simple theoretical framework that formalizes the mechanism through which generative AI could disproportionately reduce demand for junior workers and generates the testable prediction that underpins our empirical design.

In this section we adapt the canonical model of Acemoglu and Autor (2011) to our setting. The model preserves the core insight that production consists of tasks, workers specialize in different tasks, and that technology affects tasks unevenly. We depart from the original framework in two respects. First, we introduce routine intensity as the fundamental task characteristic rather than the skill of workers. Second, we model the AI shock as operating through two distinct channels, substitution and augmentation, and describe the conditions under which both channels reduce the relative demand for junior labor.

Firm output is driven by the completion of various tasks. These tasks are heterogeneous in their degree of routine intensity. Before the availability of generative AI, these tasks were divided among junior and senior workers, with junior workers concentrated in the more structured and codifiable tasks. We refer to those as AI-exposed tasks.

The increase in availability and performance of AI, modeled as a productivity shock on task-level routine intensity, alters the labor market in two channels. The substitution channel captures the displacement of junior labor in AI-exposed tasks, where AI replicates structured and codified work which is predominately supplied by junior workers. The augmentation channel captures productivity gains both junior and senior workers experience in routine tasks, wherein AI would act as a tool allowing the workers to concentrate their efforts on skills that are not yet enveloped by AI.

Together, the substitution and augmentation channels alter a firm’s optimal labor demand allocation between junior and senior workers in occupations with high AI-exposure following a shock to AI availability.

### 3.1 Model Intuition

One firm produces  $Y$  by completing tasks  $T(i)$ . The production function has two inputs, junior labor  $J$  and senior labor  $S$ , which are hired at wage  $\omega_J$  and  $\omega_S$  respectively where  $\omega_S > \omega_J$ . Within each task  $i$  the inputs are imperfect substitutes, with elasticity of substitution  $\sigma$  deciding how easily firms reallocate between them in response to relative price changes.

The aggregated production function is

$$Y = \left( \int_0^1 T(i)^{\frac{\sigma-1}{\sigma}} di \right)^{\frac{\sigma}{\sigma-1}}$$

Each task  $i$  is performed using a combination of junior labor and senior labor with a CES technology:

$$T(i) = [a_J(i)J(i)^\rho + a_S(i)S(i)^\rho]^{1/\rho}$$

Where  $a_k(i)$  captures the productivity of input  $k$  in task  $i$  and  $\eta = \frac{1}{1-\rho} > 1$  is the elasticity of substitution between inputs. This formulation of task productivity allows junior and senior labor to be used together, with their relative quantities responding to relative wages.

#### The AI Productivity Shock

AI productivity in task  $i$  is given by  $a_A(i) = \phi \times r(i)$  where  $\phi > 0$  increases after the introduction of ChatGPT in November 2022, raising the productivity of AI across all tasks, in accordance with the routine intensity  $r(i)$ . This parameter captures the level of structured codifiable work a task involves. An increase in  $\phi$  manifests itself through two channels, substitution and augmentation. The relative strength of these forces determines a firm's optimal junior-to-senior ratio in response to the shock.

#### Substitution Channel

Assumption 1: Junior workers have a comparative advantage over senior workers across tasks characterized by high routineness, meaning  $\frac{a_J(i)}{a_S(i)}$  is increasing with  $r(i)$ .

The intuition behind this is that junior roles involve structured and codified tasks, which is what is most susceptible to LLM substitution. Senior roles on the other hand require judgment, tacit knowledge, and experience accumulated over a career, which remain beyond current AI capabilities.

Comparative advantage in this model reflects differences in the opportunity cost between seniority types. It may be the case that senior workers are more productive than junior workers in all tasks  $i$ , but their higher wage makes so that allocating them to the subset of structured and codified tasks is an inefficient allocation. These tasks are thus allocated to junior workers, who have a lower opportunity cost, enabling a higher relative productivity of junior workers in these tasks.

The most routine tasks and the tasks where juniors are relatively most productive are, with this logic, one and the same. This assumption is supported by Brynjolfsson et al. (2025a),

who argue that codified knowledge, which junior workers supply relatively more of, is disproportionately substituted by AI.

By assumption 1, the tasks most exposed to AI displacement are those where junior labor is concentrated. A rise in  $\phi$  therefore reduces demand for junior labor more than senior labor. This is what constitutes the substitution channel. As AI becomes more productive in routine tasks, firms can increasingly substitute junior workers for AI services.

We maintain the substitution channel as a direct implication of Assumption 1 and the task specificity of AI productivity rather than deriving from the cost-minimization problem.

### Augmentation Channel

The AI shock raises the productivity of both junior and senior workers in AI-exposed tasks. The senior productivity:  $a_S(i) = \bar{a}_S(i) + \gamma_S \times a_A(i)$  where  $\gamma_S > 0$  captures the degree to which senior workers benefit from AI augmentation in task  $i$ .  $\bar{a}_S(i)$  is the baseline productivity of senior workers.  $\gamma_S \times a_A(i)$  captures the increase in productivity of senior workers in task  $i$  following an AI shock.

The junior productivity formulation is like that of senior workers:  $a_J(i) = \bar{a}_J(i) + \gamma_J \times a_A(i)$ , where  $\gamma_J > 0$  captures the degree to which junior workers benefit from AI augmentation in task  $i$ .

Assumption 2: For tasks with sufficiently high routine intensity,  $\gamma_S > \gamma_J$ . Meaning that the augmentation gain is higher for senior workers than or junior workers. When AI handles routine sub-tasks, workers are freed to reallocate their effort toward judgement intensive work, shifting the task continuum. Senior workers already possess the experience to perform this work productively. Junior workers, by contrast, are still in the process of accumulating these capabilities. Augmentation is therefore less valuable for juniors, because there is less experience and judgement for the technology to augment.

The findings in the Brynjolfsson et al. (2025b) call center study appears contradictory to Assumption 2. However, the generalization of the findings over all professions is questionable as the scope of judgment-intensive reallocation by seniors is limited.

### The Firm's Cost-Minimization Problem

The firm chooses different quantities of  $J(i)$  and  $S(i)$  to minimize total cost:

$$\min_{J(i), S(i)} \int_0^1 [\omega_J J(i) + \omega_S S(i)] di$$

All with regards to the CES task production function above and the aggregated output target  $Y \geq \bar{Y}$ .

## 3.2 Comparative Static

Cost minimization with CES technology yields the following relative input demand for junior versus senior labor in task  $i$ :

$$\frac{J(i)}{S(i)} = \left( \frac{a_J(i)}{a_S(i)} \right)^\eta \left( \frac{w_S}{w_J} \right)^\eta$$

Taking the logs and differentiating with respect to the AI productivity parameter  $a_A(i)$  while holding wages fixed:

$$\frac{\partial \ln[J(i)/S(i)]}{\partial \phi} = \eta \times r(i) \times \left[ \frac{\gamma_J}{\bar{a}_J(i) + \gamma_J \times \phi \times r(i)} - \frac{\gamma_S}{\bar{a}_S(i) + \gamma_S \times \phi \times r(i)} \right]$$

This is the main comparative static of our model. The sign and magnitude of the expression depend on the relation between junior and senior augmentation. The expression in the brackets captures the net augmentation and it is negative if and only if:

$$\frac{\gamma_S}{\bar{a}_S(i)} > \frac{\gamma_J}{\bar{a}_J(i)}$$

Under Assumption 1,  $\frac{a_J(i)}{a_S(i)}$  is strictly increasing in  $r(i)$ . The right-hand side of the condition therefore shrinks as tasks become more routine. Under Assumption 1 and 2, the comparative static is unambiguously negative given task  $i$  is sufficiently routine. For low routine tasks, the sign is ambiguous and the model makes no strong prediction. The magnitude of the expression is increasing with  $r(i)$ . This is what we exploit in our empirical design, comparing the evolution of junior share across occupations that differ in their baseline AI-exposure.

### 3.3 Combined Effect

The comparative static in Section 3.2 describes the augmentation channel. Under assumption 1 and 2, an increase in AI productivity lowers the relative demand for junior labor in sufficiently routine tasks as experienced workers benefit more from augmentation. The substitution force enters separately under Assumption 1. As AI becomes more productive in the routine tasks where junior workers have a comparative advantage, firms can increasingly replace junior labor with AI directly.

Combining the substitution and the net negative augmentation effect gives the paper its central testable prediction: Following the release of ChatGPT, occupations with higher AI exposure should experience a larger decline in the share of junior labor demand relative to occupations with lower AI exposure. This prediction is taken to the data in Section 4, where AI exposure indices act as proxies for the routine intensity parameter  $r(i)$ .

## 4 Data

Testing whether the public availability of generative AI disproportionately reduced the demand for junior workers requires data that simultaneously captures vacancies at the job-posting level, contains the seniority level of each role, and measures the AI exposure of the underlying occupation. We construct a novel dataset combining three elements: job postings from the Swedish Public Employment Service's JobTech Links dataset, occupational AI

exposure indices from Engberg et al. (2023) and Eisfeldt et al. (2023) and a keyword-based seniority classification applied to job posting titles. Each component is described in turn.

## 4.1 JobTech Links

The source for the job postings analyzed in this paper is a dataset called JobTech Links which is a joint project between the Swedish Public Employment Service and Sweden’s largest job boards, such as Blocket Jobb and Monster. JobTech Links generates around 30 percent more unique job postings than the Swedish Public Employment Service’s job board alone. The Swedish Public Employment Service’s job posting board alone contains around ~500,000 unique job postings each year and there are around 50,000 active job postings at a given moment. The data is compiled using an API that generates one file per day, containing all open job postings on that date from all participating job boards. Data on an individual job posting contains an ID, timestamp, job title, description, hiring organization, job location, occupational code, source links, experience requirements, and parsed keywords from the posting description. The dataset is combined across days and deduplicated by matching source links. A further step of deduplication is performed using Jaccard deduplication according to the preferred method of the dataset administrators. The dataset contains approximately 3.6 million job postings between the 30<sup>th</sup> of March 2021 and the 28<sup>th</sup> of February 2026 from over 100 thousand firms.

The dataset contains no postings between March 7 and April 26 2025, creating a gap of approximately seven weeks in the observational window. Postings that remained active beyond this period are captured in the April 27 observation, meaning that not all hiring activity during the gap is lost. However, since the probability of a posting remaining active may differ systematically by seniority, with junior roles possibly turning over faster than senior roles, retaining only the surviving postings risks introducing differential selection. We therefore exclude all postings from this period entirely

## 4.2 AI Exposure Indices

To identify the degree to which an occupation’s tasks overlap with the capabilities of AI, we use two independently constructed AI exposure indices from Eisfeldt et al. (2023) and Engberg et al. (2023), differing both in methodology and scope.

Eisfeldt et al. (2023) measure occupational exposure to GenAI capabilities by employing a Large Language Model (OpenAI’s GPT 3.5 Turbo model) to classify roughly 20,000 tasks, as defined by the O\*Net database, by how well they correspond to the known capabilities of Generative AI models. Occupational exposure scores at the 6-digit SOC level are constructed by aggregating tasks exposure scores by occupational task composition. Critically, the AI exposure score constructed is based on model capabilities around the time of the release and not those of more contemporary models, reflecting the state of the technology when we examine its effects.

Engberg et al. (2024) construct the Dynamic AI Occupational Exposure (DAIOE) index by tracking AI performance on standardized performance benchmarks between 2010 and 2023. By aggregating task-level performance scores to the occupational level annually, the index

captures the exposure of each occupation over time. We use the 2023 version of the index, reflecting the technological capabilities at the time of our natural experiment.

While the two indices are highly correlated ( $r = 0.84$ ), they are still meaningful complements. First, the two indices measure conceptually distinct dimensions of AI exposure. The DAIOE captures broad accumulated exposure to AI capabilities across multiple technological domains from 2010 up to 2022, beyond that of the capabilities of LLMs. The Eisfeldt index, by contrast, captures specifically the overlap between an occupation and the capabilities of generative AI at the moment of the ChatGPT release. Second, consistent results across both indices address the concern that any findings are purely an artifact of how AI exposure is measured in an individual AI exposure index.

#### 4.2.1 Occupational Classification Crosswalk

While the DAIOE is publicly available at the SSYK-2012 level, the Eisfeldt exposure scores are available using the United States government’s Standard Occupational Classification (SOC-10) System. As the job posting data uses Statistics Sweden’s Standard for Swedish occupational classification (SSYK-12), it is necessary to map the SOC codes to their corresponding SSYK classifications. We crosswalk SOC-2010 level occupational exposure scores to SSYK-2012 by using publicly available SOC-ISCO (International Classification for Occupations) and ISCO-SSYK crosswalks and averaging exposure scores where there are many-to-many combinations. This averaging introduces measurement error into the Eisfeldt index that is absent from the DAIOE, which is discussed further in section 7.2.

### 4.3 Summary Statistics

This section contains summary statistics broadly outlining the structure of the data. The dataset contains around 3.6 million postings across 250 weeks. The share of identified senior and junior job postings amount to around one sixth of all postings with the rest inconclusive. Senior postings outnumber junior postings roughly two to one. Postings that are neither classified as junior or senior are largely excluded from the analysis, given that the treatment affects junior and senior jobs that for some reason are not classified as such as well. The construction of junior and senior classification is described in Section 5.3.

***Table 1: Job Posting Summary Statistics***

| Total Postings | Unique Firms | Unique occupations | Unique Weeks | Share Junior | Share Senior |
|----------------|--------------|--------------------|--------------|--------------|--------------|
| 3,550,951      | 102,105      | 394                | 250          | 0.056        | 0.111        |

*Notes:* This table presents summary statistics for the whole dataset. Data source: JobTech Links (2026)

Table 2 presents summary statistics for the two AI exposure indices. The two indices are highly correlated ( $r = 0.84$ ), though the remaining variation is meaningful and the indices differ both in methodology and scope.



**Table 2: AI Exposure Index Summary Statistics**

| Index    | Mean  | SD    | Min   | P25   | Median | P75   | Max   |
|----------|-------|-------|-------|-------|--------|-------|-------|
| Engberg  | 0.657 | 0.282 | 0.004 | 0.358 | 0.715  | 0.928 | 0.999 |
| Eisfeldt | 0.389 | 0.228 | 0     | 0.130 | 0.460  | 0.626 | 0.958 |

*Notes:* This table presents summary statistics on AI exposure for whole dataset. Data source: JobTech Links (2026)

Table 3 presents summary statistics on Firm Characteristics. The distribution of postings per firm is highly skewed, with a median of 2 postings and a 90th percentile of 23 over the full observational window. This indicates that a small number of large firms account for a disproportional share of hiring activity in the dataset.

**Table 3: Firm characteristics**

| Mean | SD  | P10 | Median | P90 | N       |
|------|-----|-----|--------|-----|---------|
| 34.8 | 416 | 1   | 2      | 23  | 102,105 |

*Notes:* This table presents summary statistics for firm-level postings. Data source: JobTech Links (2026)

Figure 1 graphs the evolution of the number of monthly job postings during our observational window. Consistent with other documentation of the Swedish labor market, job postings have decreased from their post-covid peak. The dataset contains no postings between March 7 and April 26 2025, creating a gap of approximately seven weeks in the observational window.

**Figure 1: Monthly job postings from 2021-03 to 2026-03**



*Notes:* This graph depicts raw and smoothed monthly job postings between March 2021 and March 2026. Data source: Authors' rendering of daily job postings from JobTech Links (2026)

## 5 Empirical Strategy

This section presents the empirical strategy and methodology through which we test whether the increased availability of generative AI disproportionately reduced the demand for junior workers. First, we define the dependent variable. We then describe the classification of occupations by seniority level. Next, we justify the ChatGPT release date as a credible natural experiment and explain the choice of observational window. We then describe the AI exposure indices used to assign treatment intensity at the occupational level. Finally, we present the complete specification of our model and assess the identifying assumptions necessary.

### 5.1 From Model to Empirical Specification

The theoretical framework in Section 3 yields a testable prediction: the public release of ChatGPT in November 2022, an increase in  $\phi$ , shifts the ratio of junior and senior labor demand downwards in more routine based tasks  $r(i)$ . This comparative static and maintained assumptions are translated into a regression coefficient through our empirical strategy.

The increase in  $\phi$  corresponds to *Post*, an indicator for observations following the release of ChatGPT. This reasoning is defended in Section 5.4. The routineness of a task,  $r(i)$ , is empirically approximated through the assortment of occupation-level AI indices used. Junior and senior labor inputs, *S* and *J*, correspond to job postings categorized by hierarchical keyword as described in Table A3. The model’s junior-senior-ratio,  $J(i)/S(i)$ , corresponds to the junior share within an occupation-month cell, formally defined in section 5.2.

The regression in Table 6 analyses whether the productivity shift in the ratio is driven by a declining junior labor input, an increase in senior labor, or both. This will provide evidence on the extent of the substitution and augmentation mechanics of the model.

### 5.2 Dependent Variable

The dependent variable is the share of junior postings among seniority-classified postings within a given occupation-half year or firm-occupation-half year cell, defined as:

$$JuniorShare_{k,t} = \frac{Junior_{k,t}}{Junior_{k,t} + Senior_{k,t}}$$

where  $Junior_{k,t}$  and  $Senior_{k,t}$  denote the number of postings classified as junior and senior respectively within occupation  $k$  in half year  $t$ , with the firm subscript  $f$  added in the firm-half year specification. It corresponds directly to  $J(i)/[J(i)+S(i)]$  in the model.

The theoretical framework in Section 3 predicts that the release of ChatGPT shifts the cost-minimizing input decision away from junior labor in high-AI occupations. If this is true, it should manifest as a within-occupation or within-firm decline in the share of junior postings concentrated in occupations with higher AI exposure following the release of ChatGPT.

While the junior share is our primary outcome variable, it cannot on its own distinguish between two theoretically distinct mechanisms, a reduction in junior postings driven by AI substitution, and an increase in senior postings driven by AI augmentation of experienced

workers. To decompose flows into its constituent parts, we additionally estimate separate specifications using the absolute level of junior postings, senior postings and total postings as outcome variables. These level specifications complement the share analysis by allowing us to connect the empirical results to the two predictions generated by the theoretical framework in Section 3.

## 5.3 Seniority

The classification of job postings by experience or seniority level is central to our empirical design. The share of junior postings within a given occupation-half year cell or firm-occupation-half year cell is only meaningful if the classification reliably captures experience requirements or the formal hierarchical position that serves as a suitable proxy for it. In this section we describe how we identify the experience level requested in each job posting from its title. The implications of this classification for the external validity of our results are discussed in Section 7.4

### 5.3.1 Identification

To classify job postings by seniority level (J and S), we parse job titles for keywords commonly associated with requirements of previous work experience. Each posting is assigned a junior dummy, a senior dummy, or neither, depending on whether its title contains a keyword from the respective lists, which are listed in Table A3 in the appendix, followed by the number of job postings identified per keyword.

Where a posting contains both a junior and senior keyword, the senior classification takes precedence as the senior keywords are stronger and more deliberate indicators of experience requirements than junior keywords. The implications of misclassification for our estimates are discussed in Section 8.2.

### 5.3.2 Keyword selection

The keywords used to classify seniority are originally sourced from a mapping of keywords by seniority from Hosseini et al (2026) and their analysis of LinkedIn profiles and resumes from the Burning Glass dataset. We adapt their keyword list to the Swedish labor market in two ways. First, we translate the original English keywords into Swedish, retaining only those translations that carry an unambiguous experience indication or hierarchical meaning in the Swedish labor market context. Second, we extend the list with Swedish-specific terms that have no direct English equivalent but are commonly used as indicators of occupational experience in Swedish job postings, such as "Ansvarig", "Erfaren", and "Nyexaminerad".

Naming conventions are not homogenous across occupations. While some sectors have well established traditions of communicating experience level through job titles, other rely on flat organizational structures, or occupational identity labels that carry no hierarchical meaning. This heterogeneity has direct implications for the validity of a keyword-based seniority classification and motivates our interpretation of the junior-senior distinction as capturing experience requirements rather than purely hierarchical position. The validity of our classification does not rest on the presence of formal organizational hierarchy, but instead on

whether a posting signals that no prior experience is required or if a candidate must have accumulated tacit occupational knowledge that is beyond the reach of current AI capabilities.

We exclude SSYK group 1, managerial occupations, across all specifications as these roles are senior by definition and lack meaningful junior counterparts against which to identify the experience margin our classification is designed to capture.

The most consequential classification decision concerns “Assistant” and its English equivalent “Assistant”, by far the most common keyword in the junior sample. In a Swedish labor market context, the word performs dual functions. In professional occupations such as “revisionsassistent” or “projektassistent”, it serves as a credible signal of an entry-level role. In care occupations, however, the word functions primarily as an occupational identity label without explicit hierarchical connotations, such as “personlig assistent” (personal care assistant) or “vårdassistent” (care assistant). Even within care occupations, meaningful variation in experience requirements exists. A posting for a personal care assistant is different from one specifying multiple years of care experience or managerial responsibility beyond the typical tasks of the occupation. This distinction is economically meaningful regardless of whether the occupation has the formal career ladder most common in white-collar occupations. Care occupations score low on both AI indices, reflecting low routine intensity following the theoretical framework, where the model predicts little substitution effect. Nonetheless, the group of care and service workers constitute a large portion of the identified sample (Table A4). We verify the robustness of our results to the exclusion of them (SSYK 5343) and the keyword “assistant” broadly in Section 6.4.2 and discuss the implications of this exclusion for the external validity of our estimates in Section 7.4.

## 5.4 Cutoff

Having defined the dependent variable and the seniority classification, we turn to the treatment variable, the release of ChatGPT, and explain why it constitutes a credible cutoff for a natural experiment, and why it constitutes a discrete increase in  $\phi$  in the model, despite the gradual nature of technological diffusion.


The use of the release data of ChatGPT as a cutoff is not obvious in and of itself.

Technological diffusion is rarely instantaneous, by 2025 far from all Swedish organizations report using AI in their operations, and full integration of generative AI into hiring decisions likely takes considerable time following the shock. Our research design does not, however, require firms to have adopted AI by the time of the cutoff. Rather, we argue that the public release of ChatGPT constitutes a discrete change in the technological frontier available to firms. Regardless of whether any specific firm had integrated it yet, our model only requires that the release of ChatGPT meaningfully changed firm expectations, and planning horizons for their hiring decisions.

Because our cutoff-date is not a typical hard cutoff, it is also pivotal to investigate the exogeneity of the treatment. If firms had knowledge of the developing abilities of the technology, and its imminent release, they could proactively change their junior hiring. To substantiate this claim, we investigate discontinuities in related trends before and after the release. In Figure 2 and 3, the mentions of “AI” on earnings calls are plotted over time, where

we can observe a sharp increase immediately following the release of ChatGPT. The effect holds true both for the S&P500 and the STOXX Europe 600, indicating that discourse around the technology was not limited to the U.S. The use of this date, or late 2022, as a reference point for similar research has been utilized in peer-reviewed articles such as Hosseini et al. (2026) and Brynjolfsson et al. (2025b)

**Figure 2: Mentions of AI in U.S. earnings calls**



Notes: Frequency of “AI” mentions in U.S. earnings calls. The dashed vertical line denotes the public launch of ChatGPT in November 2022. Source: Hosseini & Lichtinger (2026)

**Figure 3: AI mentions in Stoxx Europe 600 earnings transcripts**



Notes: Quarterly frequency of “AI” mentions in Stoxx Europe 600 earnings call transcripts. Source: Bloomberg News (2025, March 21). AI Chatter Hits Record on European Earnings Calls Across Sectors

A further validation the release of ChatGPT as an exogenous and unexpected shock can be inferred from the stock market. Eisfeldt et al. (2023) show that the release of ChatGPT had a significant effect on equity prices in the U.S., with earnings forecasts increasing more for firms with a high share of AI exposed workers. They show that an “artificial-minus-human” portfolio earned 5 percent in the two weeks following the release of ChatGPT.

## 5.5 Observational Window

The length of the observational window is a balance between establishing the case for parallel trends, allowing sufficient time for generative AI to diffuse through firm hiring decisions, while simultaneously limiting influence from long-term structural trends. Our main

specification spans approximately five years, 18 months before, and 42 months after the release of ChatGPT on November 30<sup>th</sup> 2022. The unbalanced pre- and post is the result of a data limitation with the JobTech links dataset beginning in early 2021. While a longer pre-period would be desirable, 18 months provides sufficient observations to establish parallel trends and captures at least one full seasonal hire-cycle prior to the cutoff. The extended post-period of 42 months allows us to document both immediate as well as long-run effects reflecting the cumulative diffusion and development of the technology.

## 5.6 AI Exposure

Our ideal empirical specification would be to identify firms that adopt AI, to directly compare their junior hiring behavior against non-adopting controls, thereby isolating the causal effect of AI adoption from broader trends in hiring that are unrelated to the availability of AI. Due to data limitations, this level of identification is not available to us. Instead, we employ occupation-level measures of AI exposure as the proxy for the model’s task-level routine intensity  $r(i)$ , employing the task-based scores of Engberg et al. (2024) and Eisfeldt et al. (2023). This approach captures the degree to which the tasks constituting a given occupation are suitable to LLM substitution.

This approach is well-established in the literature. Occupation-level AI exposure measures have been used by a number of recent and influential studies of AI’s consequences on the labor market, including Brynjolfsson et al. (2024), Dominski and Lee (2025), and Eckhardt and Goldschlag (2025). The validity of this method in a Swedish context is further supported by Engberg et al. (2025a), who show that AI-exposure is a meaningful predictor of AI adoption in Sweden, holding true for both small and large firms.

We employ both AI exposure indices in parallel as both are an imperfect proxy to  $r(i)$ . Consistency across estimates, suggests that we are capturing the underlying task-level exposure rather than idiosyncrasies from either measurement approach.

## 5.7 Control Variables

The empirical specification includes a set of fixed effects to control for unobserved heterogeneity across firm, occupation, and time. Occupation fixed effects control for constant differences across job codes, including task composition and baseline demand for said job. Firm fixed effects control for differences across firms, such as size, hiring and training practices. Time fixed effects capture shocks affecting the labor market, such as macroeconomic trends and seasonal swings.

## 5.8 Model Specification

To examine whether the availability of generative AI disproportionately reduced the demand for junior workers in AI-exposed occupations, we estimate two fixed effects OLS models that exploits variation at different levels of aggregation.

### 5.8.1 Occupation-Half Year FEOLS

Our baseline specification is estimated at the occupation-half year level, by constructing a panel of occupations observed on a 6-month basis and use the share of junior job postings within each occupation-half year as the outcome variable.

$$(1) \text{JuniorShare}_{k,t} = \alpha + \beta_1 \times (AI_k \times Post_t) + \theta_k + \delta_t + \varepsilon_{k,t}$$

The coefficient  $\beta_1$  is the empirical translation of the model's comparative static at high  $\bar{a}_A(i)$ . The model intuition predicts  $\beta_1 < 0$ .  $Post_t$  is equal to 1 for every half year after the release of ChatGPT on November 30<sup>th</sup> 2022. The coefficient of interest,  $\beta_1$ , captures whether occupations with higher AI exposure experienced a differential decline in their junior job posting share following the release of ChatGPT. Occupation fixed effects,  $\theta_k$ , absorbs all time-invariant differences in junior share across occupations. Half year fixed effects,  $\delta_t$ , absorb aggregate shocks common to all observations, including seasonal vacancy patterns and macroeconomic conditions. Standard errors are clustered at the occupational level to account for serial correlation within occupations and firms over the observational window. Observations are weighted by the total number of postings in each occupation-half-year.

### 5.8.2 Firm-Half Year FEOLS

Our preferred specification moves identification to the firm-occupation-half year level, comparing changes in the junior share of postings across occupations with different levels of AI exposure within the same firm.

$$(2) \text{JuniorShare}_{f,k,t} = \alpha + \beta_1 \times (Post_t \times AI_k) + \theta_k + \mu_f \times \delta_t + \varepsilon_{f,k,t}$$

The coefficient of interest captures whether occupations with higher AI exposure, within the same firm, experienced a differential decline in their junior job posting share following the release of ChatGPT. Beyond the fixed effects included in the occupation-half year specification, we also include firm-by-time fixed effects,  $\mu_f \times \delta_t$ , which will control for firm shocks that vary over time. This design addresses the main concern with the occupation-half year specification, occupation-level economic confounders that are correlated with both AI exposure and the timing of ChatGPT. Observations are weighted by the total number of postings in each firm-occupation-half-year cell. This accounts for the high skewness in posting activity documented in Section 4.3. Weighting by cell size gives more influence to firm-occupation combinations with higher posting activity, reducing the influence of noisy cells on the estimates.

### 5.8.3 Event Studies

To examine the dynamic evolution of the treatment effect and assess the validity of the parallel trends assumption, we complement the main DiD specification with event study estimates, estimated by:

$$(3) \text{JuniorShare}_{k,t} = \alpha + \sum_{\tau \neq -1} b_\tau (1\{t - t_0 = \tau\} * AI_k) + \theta_k + \delta_t + \varepsilon_{k,t}$$

$$(4) \text{JuniorShare}_{f,k,t} = \alpha + \sum_{\tau \neq -1} b_\tau (1\{t - t_0 = \tau\} * AI_k) + \theta_k + \mu_f \times \delta_t + \varepsilon_{f,k,t}$$

The coefficients  $\beta_\tau$  trace out the dynamic effect of AI exposure on the share of junior job postings, relative to the half year immediately preceding the release of ChatGPT. Beyond the components mentioned above,  $t_0$  denotes the half year of the ChatGPT release,  $\tau$  indexes half years relative to the cutoff and  $\tau = -1$  is the omitted base period, the half year immediately before the release.

The event study captures the dynamic response of the J/S ratio to the AI shock  $\phi$ . According to the model, the estimate of interest should be approximately zero for all half years before the cutoff and become negative for all periods after. The timing would reflect the speed of firms adjusting their hiring practices to the new productivity frontier.

We aggregate postings to rolling six-month periods offset by one quarter, October through March and April through September. Standard calendar half-years are systematically different in vacancy patterns across both seniority and occupation, with the first half capturing the spring trough and the second the autumn peak. Under the standard split, the release of ChatGPT in November 2022 would place the event study reference period in a seasonally low period for junior vacancies, producing an artificially low baseline. The offset split avoids this by ensuring each period spans both part of a peak and a trough. We verify robustness to the standard calendar split in Section 6.4.1.

## 5.9 Identifying Assumptions

The causal interpretation of our fixed effects specification hinges on the parallel trends assumption, which differs slightly between our two models.

For the occupation-half year specification, the assumption requires that in the absence of the ChatGPT release, the within-occupation junior share would have evolved in parallel across occupations with different levels of AI exposure conditional on our fixed effects. The occupation fixed effects absorb time-invariant differences in junior share across occupations, and the half year fixed effects absorb aggregate shocks common to all occupations in each half year.

For the firm-half year and occupation fixed effects specification, the parallel trends assumption requires that in the absence of the widespread availability of AI, the within-firm junior share would have evolved in parallel across occupations with different levels of AI exposure, conditional on firm-half year and occupation fixed effects. Rather than assuming all occupations face the same shocks, we instead exploit variation within firms, under the more relaxed assumption that firms face common shocks that affect their vacancies homogeneously across occupations with different AI exposure.

The plausibility of parallel trends is assessed visually through event study estimates alongside the static regressions throughout Section 6. Under this assumption, the interaction coefficients between AI exposure and half-year periods should be statistically indistinguishable from zero in all pre-treatment periods.



## 6 Results

This section presents the result of our empirical analysis. We begin by estimating the seniority-biased effect of generative AI using an occupation and half-year fixed effects specification, which establishes the baseline relationship between AI exposure and the junior share of seniority-classified postings before and after the release of ChatGPT. We then present the results of our preferred model using firm-by-time and occupation fixed effects, which tightens identification by instead exploiting within-firm variation. For each specification we report the main coefficient on  $Post_t * AI_k$ , and an associated event study. We continue by presenting two alternative specifications. First examining whether the effect is concentrated in a high-skill subsample of the labor market, occupations requiring higher education. Second, examining the relationship between AI exposure and senior share of all job postings. The section concludes with a series of robustness checks that investigate the sensitivity of the main results.

### 6.1 Fixed Effect OLS – Within Occupation

We estimate the first model using two different AI indices, DAIOE from Engberg et al. (2023) and the exposure index created by Eisfeldt et al. (2023) Using the DAIOE index, the estimate on the coefficient of interest,  $Post_t * AI_k$ , is -0.076, which can be compared to the estimate using the Eisfeldt index of -0.080 with both results significant at the 0.1 and 1 percent level respectively.

**Table 4: Effect of AI exposure on occupation-level junior share**


|                           | DAIOE (Engberg)      | Eisfeldt et al.     |
|---------------------------|----------------------|---------------------|
| POST $\times$ AI Exposure | -0.076***<br>(0.019) | -0.080**<br>(0.031) |
| Dependent variable:       | Junior Share         | Junior Share        |
| Num.Obs.                  | 2,817                | 2,817               |
| Adjusted R <sup>2</sup>   | 0.945                | 0.945               |
| S.E.: Clustered           | By: SSYK             | By: SSYK            |
| FE: half-year / SSYK      | Yes                  | Yes                 |

*Notes:* This table reports difference-in-differences estimates of the effect of AI exposure on occupation-half year junior share from Equation (1). Junior share is the share of junior postings out of all seniority classified postings in occupation k in half year t. Standard errors are clustered at the occupational level.

Significance is denoted by \*p < 0.05, \*\*p < 0.01 \*\*\*p < 0.001.

We complement the static estimates with an event study presented in Figure 4, estimating a separate coefficient for each observed half-year in our sample. Pre-trend coefficients are statistically indistinguishable from zero, providing some comfort for our parallel trend assumption. Post-treatment coefficients are consistently negative, but statistically indistinguishable from zero throughout the first 12 months after the release the cutoff. The effect deepens somewhat and is significantly negative in late 2023 and onward. The apparent recovery in 2025H1 coincides with the data gap discussed in Section 4, a seasonal low point for junior hiring, mechanically increasing the estimate, though we cannot attribute it exclusively to this source.

**Figure 4: Effect of AI exposure on occupation-half-year junior share**



*Notes:* This figure plots event study estimates of the effect of AI exposure on occupation-half-year junior share, using the DAIOE index. Coefficients are estimated from Equation (3). All regressions include occupation and time-fixed effects. Standard errors are clustered at the occupational level. Vertical line indicates October 1<sup>st</sup> 2022, the beginning of the offset half year in which ChatGPT was released.

## 6.2 Fixed Effects OLS – Within Firm

Table 5 presents the results of our preferred firm-level specification estimated with firm-period and occupation fixed effects. The coefficient on  $Post_t * AI_k$  is -0.057 when estimated with the DAIOE index, significant at the 1 percent level. The Eisefeldt index produces a smaller and statistically insignificant coefficient of -0.026 though the direction is consistent with the DAIOE result. To translate the coefficient into economically interpretable terms, moving from the 25<sup>th</sup> to the 75<sup>th</sup> percentile of the DAIOE index is associated with a 3 p.p. larger decline in the junior share of seniority-classified postings. Against a pre-period baseline junior share of 33.2 percent, this represents a relative decline of 9 percent in the junior intensity of job vacancies.

**Table 5. Effect of AI exposure on firm-level junior share**


|                           | DAIOE (Engberg)      | Eisefeldt et al.  |
|---------------------------|----------------------|-------------------|
| POST × AI Exposure        | -0.057***<br>(0.015) | -0.026<br>(0.017) |
| Dependent variable:       | Junior Share         | Junior share      |
| Num.Obs.                  | 111,148              | 111,148           |
| Adjusted R <sup>2</sup>   | 0.725                | 0.725             |
| S.E.: Clustered           | By: firm             | By: firm          |
| FE: firm-half-year / SSYK | Yes                  | Yes               |

*Notes:* This table reports difference-in-differences estimates of the effect of AI exposure on firm-occupation-half-year junior share. Junior share is the share of junior postings out of all seniority classified postings in occupation  $k$  in firm  $f$  in half year  $t$ . Standard errors are clustered at the firm level. Observations are weighted by total postings in the firm-occupation-half-year cell. All regressions include firm-by-time and occupation fixed effects.

Significance is denoted by \* $p < 0.05$ , \*\* $p < 0.01$  \*\*\* $p < 0.001$ .

Figure 5 presents the event study estimates for the firm-time fixed effects specification, plotting the coefficient  $\beta_\tau$  on the interaction between relative half year indicators and AI exposure for each half year relative to the ChatGPT release. Excluding our reference point, the pre-period consists of only two observations, a consequence of the panel beginning in early 2021, limiting the visual evidence for parallel trends. The first pre-period observation in 2021H2 shows a negative coefficient of approximately -0.05, while the 2022H1 observation recovers to approximately +0.05, producing a V-shape around zero rather than a systematic trend. Following the cutoff, the coefficient remains close to zero through 2023H1 before drifting gradually and consistently negative from 2023H2 onward. Estimates turn significantly negative from 2024H1, deepening progressively through 2025 and reaching approximately -0.13 at the end of the observational window. The break in the negative trend in 2025H1 falls within the data gap described in Section 4.


**Figure 5: Event study Firm-Half Year**



*Notes:* This figure plots event study estimates of the effect of AI exposure on firm-occupation-half-year junior share using the DAIOE index. Coefficients are estimated from Equation (4), with occupation and firm-by-time fixed effects. Standard errors are clustered at the firm level. Vertical line indicates October 1<sup>st</sup> 2022, the beginning of the offset half-year in which ChatGPT was released.

To decompose our results, we re-estimate the firm-occupation-half-year event study specification using the log of junior postings, senior postings and total postings as the outcome variable respectively, with one added before taking logs to handle cells with zero counts. The coefficient on  $Post_t * AI_k$  in each case reflects whether the absolute level of that posting type changed differentially in high-AI occupations following the release of ChatGPT.

**Figure 6: Firm-level junior and senior postings margin**



*Notes:* This figure plots event study estimates of the effect of AI exposure on firm-occupation-half-year junior- and senior posts respectively, using the DAIOE index. Coefficients are estimated from Equation (4), with occupation and firm-by-time fixed effects. Standard errors are clustered at the firm level. Vertical line indicates October 1<sup>st</sup> 2022, the beginning of the offset half-year in which ChatGPT was released.

Figure 6 decomposes the main result into its constituent parts, presenting event study estimates for log junior and log senior postings separately. For junior postings, the pre-period coefficients are close to zero and flat, before declining sharply and progressively from 2024 and onward, reaching approximately -0.75 log points by 2025. This corresponds to an approximately 53 percentage decrease in junior postings when going from 0 to 1 in AI exposure. For senior postings, the coefficients are positive throughout both the pre- and post-periods, although insignificantly so after the release of ChatGPT.

**Table 6: AI effect on firm-occupation-half-year log postings**

|                             | (1)                  | (2)                     | (3)                   | (4)                | (5)               |
|-----------------------------|----------------------|-------------------------|-----------------------|--------------------|-------------------|
| POST × AI Exposure          | -0.257***<br>(0.183) | -0.239***<br>(0.157)    | -0.263*<br>(0.132)    | -0.402*<br>(0.169) | -0.154<br>(0.154) |
| Dependent variable:         | Log(Total posts)     | Log(Unclassified posts) | Log(Classified posts) | Log(Junior posts)  | Log(Senior posts) |
| Num.Obs.                    | 565,849              | 498,714                 | 113,821               | 28,260             | 84,321            |
| Adjusted R <sup>2</sup>     | 0.820                | 0.819                   | 0.833                 | 0.901              | 0.795             |
| S.E.: Clustered             | By: firm             | By: firm                | By: firm              | By: firm           | By: firm          |
| FE: firm / half year / SSYK | Yes                  | Yes                     | Yes                   | Yes                | Yes               |

*Notes:* This table presents the results of the firm-occupation-half-year specification against log postings. (1) estimates the model against log total posts. (2) estimates the model against log unclassified posts. (3) estimates the model against log classified posts. (4) estimates against log junior posts. (5) estimates against log senior posts. Using DAIOE index.

Significance is denoted by \* $p < 0.05$ , \*\* $p < 0.01$  \*\*\* $p < 0.001$ .

Table 6 presents the estimated coefficient on log postings for (1) total posts, (2) unclassified posts, (3) seniority classified posts, (4) junior posts, and (5) senior posts. The estimated coefficient on total postings (1) is -0.257 log points and significant at the 1 percent level. This estimate is very similar to, although slightly larger than, the estimated coefficient on

unclassified postings. This can be compared to the estimate on senior posts, (5), -0.154, significantly smaller than the coefficient on total postings, and insignificant at any conventional significance level. The broad decline in high-AI occupations visible in the full sample does not manifest as a significant decline in senior posts. Junior posts carry the largest coefficient of -0.402 significant at the five percent level, implying a decline of approximately 33 percent.

## 6.3 Alternative Specifications

### 6.3.1 High-Skill Heterogeneity

The results thus far have established mixed results across the full sample of seniority-classified postings. However, the theoretical framework developed in Section 3 predicts that the effect should be heterogenous across the occupational distribution, concentrated in occupations where the model's assumptions are most plausibly satisfied. First, the assumption of comparative advantages in routine tasks relative to senior workers in the same occupation. This assumption holds most true in professional occupations where work differs meaningfully along the axis of seniority. In lower-skill occupations this distinction is less sharp, where junior and senior workers may differ in experience while the content of work remains similar. Second, the augmentation channel of our theoretical framework in which experienced workers can utilize their freed capacity by redirecting their time toward judgement-intensive work requires the existence of such work. This condition is most likely to hold true in high-education occupations.

We therefore restrict our sample to SSYK occupational codes that begin with 2, occupations that require advanced higher education and re-estimate our baseline specifications. Beyond the theoretical motivation, this restriction also addresses the measurement concerns discussed in Section 5. Hierarchical titling conventions are most established in high-skill professional occupations, meaning the seniority classifications are most valid and the estimated effect is least contaminated by the measurement error introduced by our seniority-classification.

**Table 7: Effect of AI exposure on junior share (SSYK 2)**

|                           | Occupation-half-year |                     | Firm-half-year         |                        |
|---------------------------|----------------------|---------------------|------------------------|------------------------|
|                           | DAIOE (Engberg)      | Eisfeldt et al.     | DAIOE (Engberg)        | Eisfeldt et al.        |
| POST $\times$ AI Exposure | -0.148***<br>(0.028) | -0.133**<br>(0.041) | -0.069*<br>(0.029)     | -0.037<br>(0.028)      |
| Dependent variable:       | Junior Share         | Junior Share        | Junior Share           | Junior Share           |
| Num.Obs.                  | 973                  | 973                 | 42,297                 | 42,297                 |
| Adjusted R <sup>2</sup>   | 0.843                | 0.842               | 0.709                  | 0.708                  |
| S.E.: Clustered           | By: SSYK             | By: SSYK            | By: firm               | By: firm               |
| FE:                       | Time / SSYK          | Time / SSYK         | Firm-by-time /<br>SSYK | Firm-by-time /<br>SSYK |


*Notes:* This table presents the results of the occupation-half-year and firm-occupation-half-year specification against junior share. Standard errors are clustered at the occupation- and firm-level respectively. Coefficients are estimated from Equation (3) and (4) and include time and occupation fixed effects and firm-by-time and occupation fixed effects respectively.

Significance is denoted by \* $p < 0.05$ , \*\* $p < 0.01$  \*\*\* $p < 0.001$ .

Table 7 presents the results from our baseline specification when restricting to only occupations requiring higher education. In the occupation-half-year specification, both estimates are negative and statistically significant, with a one unit increase in AI exposure associated with a decline in junior share of 14.8 percentage points using DAIOE and 13.3 percentage points using Eisfeldt. Moving to the firm-half-year specification, the DAIOE coefficient falls to -0.060 and remains significant at the five-percent level. The Eisfeldt coefficient falls to -0.022 and loses significance. Scaling by the interquartile range of AI exposure within SSYK group 2 occupations, a move from the 25<sup>th</sup> to the 75<sup>th</sup> percentile of AI exposure, a range of 0.158, is associated with a decline in junior share of roughly 1 percentage point. Against a baseline junior share of 18 percent in SSYK group 2 this represents a 5.6 percent decrease in junior share.


Figure 7 and 8 present event study estimates for the occupation-half-year and firm-half-year specification respectively when restricted to SSYK group 2. Both figures display flat pre-trends, supporting the parallel trends assumption underlying the difference-in-differences design. Following the release of ChatGPT, the coefficients remain near zero through 2023 before declining progressively from 2024. The break in the downwards trend coincides with the data gap mentioned in Section 4.

**Figure 7: Occupation-half-year event study (Only SSYK group 2)**



*Notes:* This figure plots event study estimates of the effect of AI exposure on occupation-half-year junior share. Coefficients are estimated from Equation (3) using the DAIOE et al. index. All regressions include occupation, and time-fixed effects. Standard errors are clustered at the occupation level. Vertical line indicates October 1<sup>st</sup> 2022, the beginning of the offset half-year in which ChatGPT was released.

**Figure 8: Firm-occupation-half-year event study (SSYK 2)**



*Notes:* This figure plots event study estimates of the effect of AI exposure on firm-occupation-half-year junior share. Coefficients are estimated from Equation (4) using the DAIOE index, with occupation and firm-time-fixed effects. Standard errors are clustered at the firm level. Vertical line indicates October 1<sup>st</sup> 2022, the beginning of the offset half-year in which ChatGPT was released.

### 6.3.2 Senior Share

To make fuller use of our dataset, we re-estimate our baseline firm-level specification using senior share of total postings as the outcome variable. Table 8 presents the estimated coefficients on  $Post_t * AI_k$  using the DAIOE and Eisfeldt et al. index respectively. The coefficients are statistically indistinguishable from zero across both specifications.


**Table 8: AI effect on firm-occupation-half-year senior share**

|                           | DAIOE (Engberg)   | Eisfeldt et al.   |
|---------------------------|-------------------|-------------------|
| POST $\times$ AI Exposure | -0.001<br>(0.005) | -0.010<br>(0.007) |
| Dependent variable:       | Senior Share      | Junior share      |
| Num.Obs.                  | 563,321           | 563,321           |
| Adjusted R <sup>2</sup>   | 0.642             | 0.429             |
| S.E.: Clustered           | By: firm          | By: firm          |
| FE: firm-half-year / SSYK | Yes               | Yes               |

*Notes:* This table reports difference-in-differences estimates of the effect of AI exposure on firm-half year senior share. Senior share is the share of senior postings out of all postings in occupation  $k$  in firm  $f$  in half year  $t$ . Standard errors are clustered at the firm level. Observations are weighted by total postings in the firm-occupation-half-year cell. All regressions include firm-by-time and occupation fixed effects. Significance is denoted by \* $p < 0.05$ , \*\* $p < 0.01$  \*\*\* $p < 0.001$ .

The event study in Figure 9 illustrates the dynamic evolution of the senior share effect. The leftmost pre-period observation is elevated at approximately 0.04 before declining toward zero at the cutoff. Following the release of ChatGPT, coefficients remain close to zero through 2023 and 2024 before slightly shifting upward in the far right of the panel.

**Figure 9: Firm-occupation-half-year event study (Senior share)**



*Notes:* This figure plots event study estimates of the effect of AI exposure on firm-occupation-half-year senior share. Coefficients are estimated from Equation (4) using the DAIOE index, with occupation and firm-time-fixed effects. Standard errors are clustered at the firm level. Vertical line indicates October 1<sup>st</sup> 2022, the beginning of the offset half-year in which ChatGPT was released.

## 6.4 Robustness Checks

### 6.4.1 Seasonality

As discussed in Section 5, junior vacancies in high-AI occupations exhibit a strong seasonal pattern, with the first half of the calendar year associated with a peak and the second half with a trough. Our preferred specification offsets the half-year boundaries by one quarter, absorbing part of this variation. Table 9 presents the static difference-in-differences estimates on the  $Post_t \times AI_k$  term using the standard calendar half-year split as a robustness check. The occupation-half-year DAIOE coefficient remains statistically significant, although the estimate is attenuated from -0.076 in the main specification to -0.038. The remaining specifications produce insignificant estimates with mixed signs, with the Eisfeldt et al. generating a positive coefficient in the firm-by-time specification.

Figures 10 and 11 present the corresponding event studies. In the firm-occupation-half-year, Figure 10, pre-period coefficients oscillate around zero, exhibiting no clear upwards or downwards trend, reflecting the seasonal reference period problem that motivates our offset split. The reference period at the cutoff falls in a seasonally low period for junior vacancies in high-AI occupations, making autumn observations appear elevated and positive by construction. During the post-period, the seasonal pattern continues during 2023 before following a qualitatively similar trajectory to our preferred specification. The seasonal oscillation reduces precision but does not alter the direction or timing of the effect when using the DAIOE index. Figure 11 presents the occupation-half-year event study under the standard calendar half-year specification, illustrating a very similar pattern to Figure 9. Despite the noise, focusing on the seasonally equivalent periods, a pattern of both lower peaks and lower troughs emerges. First half-year troughs that recover close to zero during the




pre-period and through 2023 progressively shift downward from 2024 onward, with the same holding true for the respective peaks during the second half of the year.

**Table 9: Firm-by-time calendar half-year specification**

|                           | (1)                |                   | (2)                      |                          |
|---------------------------|--------------------|-------------------|--------------------------|--------------------------|
|                           | DAIOE<br>(Engberg) | Eisfeldt et al.   | DAIOE<br>(Engberg)       | Eisfeldt et al.          |
| POST $\times$ AI Exposure | -0.038*<br>(0.015) | -0.036<br>(0.024) | -0.013<br>(0.015)        | 0.016<br>(0.016)         |
| Dependent variable:       | Junior Share       | Junior Share      | Junior Share             | Junior share             |
| Num.Obs.                  | 2,965              | 2,965             | 108,943                  | 108,943                  |
| Adjusted R <sup>2</sup>   | 0.930              | 0.930             | 0.647                    | 0.646                    |
| S.E.: Clustered           | By: SSYK           | By: SSYK          | By: firm                 | By: firm                 |
| FE:                       | Half-year / SSYK   | Half-year / SSYK  | Firm-half-year /<br>SSYK | Firm-half-year /<br>SSYK |


*Notes:* This table reports difference-in-differences estimates of the effect of AI exposure on (1), occupation-half-year junior share and (2), firm-occupation-half-year junior share. Observations are weighted by total postings in the occupation-half-year, firm-occupation-half-year cell respectively. Significance is denoted by \* $p < 0.05$ , \*\* $p < 0.01$  \*\*\* $p < 0.001$ .

**Figure 10: Calendar firm-occupation-half-year event study**



*Notes:* This figure plots event study estimates of the effect of AI exposure on firm-occupation-half-year junior share using the DAIOE index. Coefficients are estimated from Equation (4) using the DAIOE index, with occupation and firm-by-time fixed effects. Standard errors are clustered at the firm level. Vertical line indicates 1<sup>st</sup> July 2022, the beginning of the half-year in which ChatGPT was released.

**Figure 11: Calendar occupation-half-year event study**



*Notes:* This figure plots event study estimates of the effect of AI exposure on occupation-half-year junior share using the DAIOE index. Coefficients are estimated from Equation (3), with occupation-half-year fixed effects. Standard errors are clustered at the occupation level. Vertical line indicates 1<sup>st</sup> July 2022, the beginning of the half-year in which ChatGPT was released.

## 6.4.2 Exclusions

Table 10 assesses the sensitivity of the baseline estimate to the exclusion of the two most prevalent occupations in the seniority-classified sample. Column 1 estimates the baseline firm-occupation-half-year estimate of -0.057. Column 2 excludes personal care assistants (SSYK 5343), the most common occupation in the junior-classified dataset.. The coefficient attenuates modestly to -0.052 but remains significant at the five percent level. Column 3 excludes software engineers (SSYK 2512), an occupation with simultaneously very high AI exposure (DAIOE of 0.989) and high prevalence across both the senior and junior sample. The coefficient attenuates to -0.042 but remains significant at the 5 percent level. Column 4 excludes all vacancies with the keyword assistant, the most prevalent junior keyword. The coefficient attenuates slightly to -0.05 but remains significant, albeit instead at the 1 percent level.

**Table 10: Exclusions – Firm-level**

|                             | (1)                  | (2)                 | (3)                | (4)                 |
|-----------------------------|----------------------|---------------------|--------------------|---------------------|
| POST × AI Exposure          | -0.057***<br>(0.015) | -0.052**<br>(0.016) | -0.042*<br>(0.018) | -0.050**<br>(0.015) |
| Dependent variable:         | Junior Share         | Junior Share        | Junior Share       | Junior Share        |
| Num.Obs.                    | 111,148              | 108,084             | 98,777             | 106,591             |
| Adjusted R <sup>2</sup>     | 0.725                | 0.596               | 0.714              | 0.691               |
| S.E.: Clustered             | By: firm             | By: firm            | By: firm           | By: firm            |
| FE: firm / half-year / SSYK | Yes                  | Yes                 | Yes                | Yes                 |

*Notes:* This table presents the results of the firm-occupation-half-year specification with junior share as the outcome variable. (1) is the baseline specification using the DAIOE index. (2) excludes personal care assistants (5343). (3) excludes software engineers (2512). (4) excludes the keyword “assistant” completely. Significance is denoted by \*p < 0.05, \*\*p < 0.01 \*\*\*p < 0.001.

## 7 Discussion

### 7.1 Interpretation of Findings


The occupation-half-year specification establishes a baseline negative relationship between AI exposure and junior share following the release of ChatGPT, significant under both indices. The effect does not emerge immediately but deepens gradually toward the end of the observational window, consistent with AI adoption operating through a diffusion process rather than firms responding instantaneously to the availability of a new technology. While the occupation-half-year specification envelops all our classified observations, it is more exposed to confounding from macroeconomic shocks that differentially affect high-AI and low-AI sectors and should be interpreted with this in mind.

Moving to our preferred firm-half-year specification, identification tightens considerably by absorbing all firm-level time-varying shocks, meaning the estimate is no longer confounded by differential trends across industries. The DAIOE estimate remains negative and significant while the Eisfeldt estimate attenuates and loses significance. Individual estimates are estimated with a high degree of imprecision given the coarser half-year aggregation and demanding firm-by-time fixed effects structure, the overall pattern is consistent with a gradual and persistent decline in junior share in high-AI occupations following the availability of generative AI, with the effect accumulating over time rather than occurring as a discrete break at the cutoff.

This within-firm design is more demanding in two respects. First, for a firm to contribute to identification, it must hire across multiple occupations with different AI exposure scores. Consider a firm whose vacancies are concentrated in high-AI-exposure occupations, such as software engineering. Even if AI adoption causes junior postings in that occupation to decline relative to senior postings, this variation is entirely absorbed by the firm-half-year fixed effects, which already conditions out any common shift in junior vacancies across that firm's postings meaning that the observation does not contribute to our identification. Second, even for firms that hire across a diverse set of occupations, identification requires that the cross-occupational variation in AI exposure is sufficient to lead to differential trends in junior vacancies. The within-firm variation between high- and low-exposure occupations may be too small to produce detectable differential effects. If a firm were to respond to the AI shock by reducing junior vacancies broadly across all their posted occupations rather than selectively, and in proportion to AI exposure, the firm-occupation-half-year specification absorbs this.

Figure 12 presents a histogram of within-firm spread in AI exposure, defined as the difference between the maximum and minimum AI exposure across occupations posted by the same firm, excluding firms that post only in one occupation. The distribution spans the full range from zero to one, confirming that a substantial set of firms hire across occupations with heterogeneous AI exposure. The concentration of firms at low spread values is consistent with the attenuation observed in the firm-half-year specification relative to the occupation-half-year results, as many firms hire within a narrow range of the AI exposure distribution, contributing little identifying variation to the within-firm comparison.

**Figure 12: Within-firm max difference in AI exposure**



*Notes:* This figure plots the prevalence of differences within firms between the highest AI exposure posting and lowest AI exposure posting, excluding firms that only hire 1 occupation. Data source: Authors' rendering of within-firm max difference in AI exposure from JobTech Links (2026)

Two spikes in the distribution warrant specific attention. The spike visible at approximately 0.25 is driven by 806 firms hiring a combination of software developers (2512) and commercial sales representatives (3322), two occupations that appear together frequently and sit at different points of the AI exposure distribution. The smaller spike around 0.06 is driven by 235 firms whose largest internal difference is between advertising and marketing professionals (2431) and software developers (2512). Notably, software developers appear as the high-exposure reference point in two of the most common pairings, meaning that the exclusion of software engineers reported in Table 10 removes a significant share of the within-firm identifying variation for many firms. The fact that the coefficient from our exclusion robustness check remains significant despite the exclusion of software engineers suggests that other combinations with AI exposure differences also contribute to the estimate.

The decomposition by posting type reveals the mechanism underlying the junior share decline. The contraction in our seniority-classified sample is driven by fewer junior postings rather than more senior postings, consistent with the substitution channel in our theoretical framework, where AI capabilities reduce the optimal quantity of junior labor while leaving senior labor demand largely intact. This decomposition indicates that the overall contraction in high-AI occupations since the hiring peak of 2022 is not uniform across seniority, instead concentrated on the junior and unclassified margin. Under the interpretation that unclassified postings represent a sort of mid-level between junior and senior, which should be cautioned, one can also see a pattern of attenuating results, with the largest negative effect found in junior posts and attenuating slightly to unclassified posts and completely for senior posts. This gradient maps directly onto the theoretical prediction that AI substitution is strongest where routine intensity is highest, at the bottom the seniority distribution, and weakest where judgement and tacit knowledge dominate, at the top. Nonetheless, the insignificant coefficient on senior posts is itself informative. Against a backdrop of broad decline across all other

posting categories in high-AI occupations, seniors posts are the sole category showing no significant relationship with AI exposure.

Restricting the sample to high-skill professional occupations strengthens the results considerably, consistent with the theoretical assumptions being most plausible in this segment, with both the comparative advantage of junior workers in routine tasks and the augmentation channel for senior workers sharpest where the character of work differs more along the seniority axis. The attenuation between the specifications is large, with over half of the occupation-half-year effect disappearing once controlling for firm-by-time shocks, suggesting that part of the occupation-half-year estimate is driven by firm-level confounders correlated with AI.

Finally, the senior share specification provides a complementary view from the top of the seniority distribution. Senior postings are the most clearly and consistently classified in our dataset, with senior keywords carrying unambiguous signals of experience requirements across a broad range of occupations and industries. Using total postings as the denominator allows us to exploit the full dataset rather than only the classified subsample. If AI is reshaping the seniority composition of hiring from the bottom up, senior share should increase in high-AI occupations as junior and unclassified posts contract beneath a stable senior hiring level. The absence of a significant increase is consistent with an imprecisely estimated coefficient on senior hiring and a relatively small share of senior posts compared to unclassified posts.

The calendar half-year robustness check confirms that the direction and timing of the AI effect are not sensitive to the choice of time aggregation. The post-treatment trajectories follow a qualitatively similar pattern across both specifications under standard calendar split. Effect sizes are attenuated across all specifications relative to the offset split, a result of the reference period landing in a seasonal low with which all other estimates are compared. The calendar half-year results are presented as a robustness check, noting that the direction and timing of the effect are not sensitive to the choice of time aggregation, while statistical significance and the smoothness of the trajectory are.

Taken together, we find mixed results across specifications and outcome variables. The occupation-half-year specification has the largest estimate, with an attenuation across both AI indices when introducing the firm-occupation-half-year specification, where only the DAIOE index generates significant results. The decomposition by posting type reveals an asymmetry between seniors and all other groups in relationship with AI exposure, finding no significant relationship between AI exposure and senior posts, while other groups have significantly negative relationships. The high-skill subsample strengthens the results, with the cleanest pre-trends and the strongest theoretical motivation. The overall pattern is consistent with generative AI disproportionately displacing junior hiring in high-exposure occupations, while leaving senior demand relatively intact. However, potential confounders exist that our empirical method cannot account for, further discussed in Section 7.3, warranting caution in drawing strong causal conclusions.

## 7.2 AI Index Heterogeneity

The results from Section 6 tell the story of notable variation across the two AI indices used. The DAIOE index consistently yields larger estimates than the Eisfeldt index. Understanding why the indices produce different results requires an understanding of each measure and how they differ methodologically.

The DAIOE index, Engberg et al. (2024), captures the accumulated exposure of occupations to AI capabilities between 2010 and 2023. It is constructed by benchmarking AI performance over time, reflecting the frontier of automation potential by occupation. The Eisfeldt index is narrower in scope and more temporally precise. It measures the overlap between an occupation's tasks and the capabilities of generative AI. This difference makes the Eisfeldt index a better instrument for our empirical design as it is more in line with our theoretical framework's design to capture the exact technology whose introduction we exploit as our natural experiment.

Despite the theoretical advantages of the Eisfeldt index, it performs poorer empirically. A potential explanation to this would be that the Eisfeldt index is constructed using U.S. occupational classifications (SOC) and must therefore be cross walked to Swedish SSYK indirectly through ISCO codes. A second reason could be that LLM capabilities have continually improved since the creation of the index. If occupations that were moderately exposed in 2022 become highly exposed by 2024 or 2025, the Eisfeldt index would understate their true exposure at the time the effect materializes in our data. While this is true for our use of the DAIOE index as well, its broader scope and exposure to other technological domains may better approximate the relative exposure of occupations by AI exposure, even if the absolute levels are outdated at the point when we observe our results.

Lastly, the correlation between indices is high ( $r = 0.84$ ), but the remaining variation could be concentrated in occupations that contribute the most to the identification in the firm-half-year specification. This would explain the attenuation observed in the Eisfeldt estimates. However, the pattern of consistency of sign across both indices provides support to the idea that our findings are not driven by any idiosyncrasies of either index.

## 7.3 Limitations

### 7.3.1 Sample Bias

The identification, and differentiation, of junior and senior roles is pivotal to the empirical design of this thesis. Approximately one sixth of all postings are flagged as either junior or senior. This is a significant reduction in the size of the dataset and could give rise to biases if the identification of seniority is not conditionally random. The central concern is that the use of explicit seniority markers in job titles is not uniformly distributed across the labor market. Certain industries, notably finance, consulting, law and technology, have well-established conventions of using hierarchical titles such as "junior analyst" or "senior associate", making their postings systematically more likely to be captured by our identification method. Industries most likely to use hierarchical titles, such as finance, engineering, or law, are also

among those scoring highest on our AI exposure indices. This creates a risk that the estimated effect partly reflects sample composition rather than a labor market-wide phenomenon.

More broadly, the estimated coefficient is best interpreted as the effect for occupations where hierarchical titling conventions are most common, and generalization to the wider labor market requires the assumption that the effect would be similar in occupations where experience requirement is not explicitly signaled in job titles. It is plausible that the occupations where firms do not use explicit seniority titles are also those where distinction in tasks content between juniors and seniors is less clear. In this case, the theoretical framework underlying the work would predict that the restriction to classified posts does not miss a broader effect, instead capturing the part of the labor market where both the junior-senior divide is the sharpest and where AI driven compositional shifts are most likely to occur.

To investigate to what extent sample bias may affect the interpretation of our results, we compare the AI exposure characteristics of the seniority-classified and unclassified portions of the dataset. Table 11 below presents the average DAIOE AI-exposure score for junior job postings, senior job postings and a comparison between the seniority-classified dataset and the unclassified dataset. We see that the average AI exposure for the classified dataset is approximately 0.66, which can be compared the average AI exposure of 0.56 for the job postings that have been identified as neither junior nor senior. The difference of 0.1 suggests that the hierarchical naming of job postings does not occur uniformly across the AI exposure distribution, with postings that carry explicit seniority markers in their titles tending to come from occupations that are more exposed to AI than those that do not. That said, the magnitude of the difference, 0.1 points on a scale from 0 to 1, is modest, and the two distributions overlap substantially.

**Table 11: AI exposure across seniority and classification status**

| Panel                            | Group          | Mean AI | SD AI | N         |
|----------------------------------|----------------|---------|-------|-----------|
| A: Junior vs Senior (identified) | Junior         | 0.536   | 0.263 | 199,950   |
| A: Junior vs Senior (identified) | Senior         | 0.729   | 0.267 | 394,056   |
| B: Identified vs Non-identified  | Identified     | 0.657   | 0.282 | 594,006   |
| B: Identified vs Non-identified  | Not identified | 0.555   | 0.277 | 2,956,945 |

*Notes:* This table presents the mean and standard deviation in DAIOE AI exposure for Junior and Senior posts respectively, as well as for our seniority-classified dataset and non-classified dataset. Data source: JobTech Links (2026)

### 7.3.2 Job Title Relabeling

A potential concern with our seniority classification is the relabeling of job titles in a way that is correlated with AI exposure. Two distinct forms of relabeling could bias our estimates.

The first is upward labelling. If firms increasingly relabel junior or mid-level roles with more senior-sounding titles in a way that is correlated AI exposure, postings that do not inhabit the judgement and tacit knowledge based occupational task content of our theoretical framework

would be captured by our senior keywords rather than in the unclassified sample or potentially the junior sample. This would mechanically reduce the junior share and increase measured senior share in high-AI occupations. If such a process has taken place gradually over the length of our panel, this would inflate our estimated coefficient. The stability of senior postings in high-AI occupations visible in table 5 provides some evidence against this concern. If upward labelling were driving this result, we would expect senior postings to increase in high-AI occupations following the cutoff, which it does not. However, it does not completely rule it out since AI-driven reductions in senior demand could obscure a relabeling-driven increase.

The second is the removal of explicit seniority marker entirely, for example, “Junior analyst” becoming simply “Analyst”. This would shift postings from the junior sample to the unclassified sample. Unclassified posts decline at a slightly lower rate than total postings in high-AI occupations, which is weakly consistent with a compensating inflow of relabeled junior vacancies partially offsetting the broader decline in high-AI hiring. However, this pattern is also consistent with unclassified posts being predominantly mid-level roles that are less exposed to AI substitution than junior posts and therefore declining less sharply. We cannot distinguish between these two interpretations with the available data, and it cannot be fully ruled out as a contributing factor to our estimated decline in junior share.

### 7.3.3 Economic Confounders

The 2022-2024 period was subject to a substantial monetary tightening cycle in Sweden with the release of ChatGPT in November 2022 coinciding with the Central Bank raising the rate from zero in early 2022 to 1.75 percent in November 2022 and 4 percent by October 2023, Sveriges Riksbank (n.d). This could have disproportionately affected job posting practices in a manner that is correlated with AI exposure.

The aggregate decline in total postings visible in Table 5 is consistent with a broad contraction in hiring activity in high-AI-exposure occupations, which monetary tightening could explain. However, this interpretation weakens when considering the asymmetric development of junior and senior postings. Senior postings show no significant relationship with AI-exposure while junior posts fall most sharply of all categories. In an economic contraction we would expect both junior and senior job vacancy posts to decrease, although perhaps not at the same level. The absence of any senior decline is difficult to fully reconcile with a monetary tightening explanation, even when accounting for junior vacancy posts being more cyclically exposed. This is further supported by Engberg et al. (2026), who exploit the sharp rate increases of 2022 to test whether AI-exposed occupations reacted differently in their job posting patterns, finding no significant relationship between AI-exposure and rate driven hiring fluctuations.

Second, the firm-half year specification accounts for all firm-level time-varying shocks, including differential exposure to interest rates. For monetary tightening to confound the firm-half-year estimate it would need to differentially affect junior relative to senior vacancies within high-AI occupations compared to low-AI occupations, within the same firm and period.



Third, event studies show no significant decline in junior postings immediately following the sharp rate increases in 2022 and early 2023, with the effect only emerging from 2024 onward, consistent with a gradual diffusion of AI technologies rather than a purely economic explanation. However, this could be partially explained by a lag between the monetary tightening and its labor market effects, meaning that the timing of the effects is not conclusively evidence against a macro-driven explanation.

Although the economic confounding scenarios becomes increasingly specific, it cannot be fully ruled out. Junior vacancies in high-AI-exposure occupations may be differentially sensitive to economic cycles relative to junior posts in low-AI-exposure occupations. This remains a limitation that our design cannot fully account for.

## 7.4 External validity

Factors limiting the external validity of our analysis stem from both the method of seniority classification and the underlying data. Since we classify seniority using job title keywords, the classification only captures postings with explicit hierarchical keywords, only one sixth of all postings. This means that the estimated effect is best interpreted as local to occupations using a hierarchical titling convention, rather than labor marker-wide occurrence. The robustness tests in Section 6.3 on high-skill heterogeneity indicates the effect is the strongest and parallel pre-trends most satisfied in this segment, but not significant across all AI indices.

## 7.5 Comparison to Other Studies

Our closest point of comparison is the contemporaneous working paper by Engberg et al. (2026), made available in late March 2026. While their approach uses worker age rather than job title keywords, it provides a useful benchmark against which to assess our findings. They report that employment for 22–25-year-olds in high-AI occupations decreased by -5.5 percent relative to less exposed occupations by early 2025, with the effect first emerging during the first half of 2024. When isolating 22-25-year-olds in top-quartile AI-exposed occupations they find that employment for this group fell by 44 percent from October 2022 to early 2025, providing useful context for our relatively large estimates in the SSYK group 2 subsample when employing the same AI exposure index.

Although significance across our specifications varies, the general pattern in both our firm-by-time fixed effects model for our full sample and SSYK group 2 subsample is consistent with the effect first materializing during 2024. Considering magnitude, the estimate on our baseline firm-by-time fixed effects model is equivalent to a 9 percent decline in the junior share of postings, compared to their finding of -5.5 percent for 22-25-year-olds relative to older workers within the same employer. Our outcome variable differs from theirs, which limits direct comparability. Our approach of classifying seniority by keywords likely overlooks many postings for entry-level positions that are not explicitly labeled as such, meaning our measure captures a narrower subset of the most obviously junior roles. A consequence of this may, our estimates may be biased upward relative to theirs.

Despite this difference, the direction and timing of our effects are broadly consistent with theirs. The magnitude of the effects is also informative, although it should be interpreted with caution given the differences in outcome definition and sample composition.

## 8 Conclusion

This thesis has investigated whether the widespread availability of Generative AI has disproportionately reduced the relative demand for junior workers in the Swedish labor market. Using approximately 3.6 million job postings from Sweden’s largest job boards spanning 2021 to 2026, we classify posting seniority from job titles using a novel keyword-based approach and estimate a fixed-effects difference-in-differences design that exploits variation in occupational AI exposure around the release of ChatGPT in November 2022.

We make two contributions. First, we introduce job posting titles as a novel measure of seniority, capturing the experience firms demand at the point of hiring. Second, we contribute evidence from the Swedish institutional context to an emerging and contested literature on whether generative AI is reshaping entry-level hiring.

Across both specifications, we find negative point estimates on the interaction between AI exposure and the post-ChatGPT period. In our preferred firm-half-year specification, which identifies the effect from differential changes in the seniority composition of vacancies across occupations within the same firm, a move from the 25<sup>th</sup> to the 75<sup>th</sup> percentile of the DAIOE AI exposure index is associated with a 9 percent relative decline in the junior share of seniority-classified postings, significant at the one percent level.

Across all specifications, the effect is gradual rather than immediate, emerging clearly from 2024 onward, consistent with AI adoption operating through a diffusion process rather than a discrete shock at the cutoff. Decomposing our result by absolute postings we find that the decline in junior share is driven by a contraction in junior postings rather than an expansion of senior postings, supporting the substitution channel of our theoretical framework while also arguing against a purely macroeconomic explanation. Notably, we find no significant relationship between AI exposure and senior postings, indicating that the broad decline in high-AI occupations is concentrated further down the seniority distribution. The effect is most precisely estimated, provides the cleanest pre-cutoff parallel trends, and is largest in magnitude when restricting to high-skill professional occupations (SSYK group 2), where both the theoretical assumptions and the validity of our seniority classification are most plausibly satisfied.

These findings should be interpreted with several limitations in mind. The keyword-based seniority classification captures only the subset of postings where firms explicitly signal experience requirements in job titles, which is more common in high-skill professional occupations and in sectors with established hierarchical titling conventions. Additionally, the 18-month pre-period limits the visual evidence for parallel trends, particularly in our main specification. Finally, the possibility of differential sensitivity of junior vacancies correlated with AI exposure cannot be fully ruled out.

Our findings provide some evidence of generative AI reshaping the entry points into the labor market, reducing firms' demand for junior workers performing routine and codifiable tasks that AI is best suited to replace. The effect is gradual, meaning that our observed effects are merely an early sign of more fundamental labor market restructuring with the technology continuing to develop. To get a fuller understanding of this phenomenon we suggest continued research using firm-level AI adoption rather than AI indices as a proxy, to more precisely identify causal channels and trace whether decreasing entry-level positions have lasting long-run effects.

## 9 Bibliography

- Acemoglu, D., & Johnson, S. (2024). Learning from Ricardo and Thompson: Machinery and labor in the early industrial revolution and in the age of artificial intelligence. *Annual Review of Economics*, 16, 597–621. <https://doi.org/10.1146/annurev-economics-092123-111706>
- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30. <https://doi.org/10.1257/jep.33.2.3>
- Acemoglu, D., & Restrepo, P. (2020). Robots and jobs: Evidence from U.S. labor markets. *Journal of Political Economy*, 128(6), 2188–2244. <https://doi.org/10.1086/705716>
- Autor, D. H. (2024). *Applying AI to rebuild middle class jobs* (NBER Working Paper No. 32140). National Bureau of Economic Research. <https://doi.org/10.3386/w32140>
- Autor, D. H., Levy, F., & Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *Quarterly Journal of Economics*, 118(4), 1279–1333. <https://doi.org/10.1162/003355303322552801>
- Bick, A., Blandin, A., Deming, D. J., Fuchs-Schündeln, N., & Jessen, J. (2026). *Mind the gap: AI adoption in Europe and the U.S.* (NBER Working Paper No. 34995). National Bureau of Economic Research. <https://doi.org/10.3386/w34995>
- Bloomberg News. (2025, March 21). AI chatter hits record on European earnings calls across sectors. *Bloomberg*. <https://www.bloomberg.com/news/articles/2025-03-21/ai-chatter-hits-record-on-european-earnings-calls-across-sectors>
- Brynjolfsson, E., Chandar, B., & Chen, R. (2025a). *Canaries in the coal mine? Six facts about the recent employment effects of artificial intelligence*. Stanford Digital Economy Lab. <https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/>
- Brynjolfsson, E., Li, D., & Raymond, L. (2025b). Generative AI at work. *Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>
- Bureau of Labor Statistics. (2015). *Standard Occupational Classification Policy Committee documentation*. U.S. Bureau of Labor Statistics.
- Chandar, B. (2025). *Tracking employment changes in AI-exposed jobs* (SSRN Working Paper). Social Science Research Network. <https://doi.org/10.2139/ssrn.5384519>
- Cronert, A. (2019). *Is regulatory compliance by employers possible without enforcement? Evidence from the Swedish labor market* (Working Paper 2019:23). Institute for Evaluation of Labour Market and Education Policy (IFAU).

- Dillender, M., & Forsythe, E. (2022). *Computerization of white collar jobs* (NBER Working Paper No. 29866). National Bureau of Economic Research. <https://doi.org/10.3386/w29866>
- Dominski, J., & Lee, Y. (2025). *Advancing AI capabilities and evolving labor outcomes* (arXiv preprint No. 2507.08244). <https://arxiv.org/abs/2507.08244>
- Eckhardt, S., & Goldschlag, N. (2025, August 10). *AI and jobs: The final word (until the next one)*. Economic Innovation Group.
- Eisfeldt, A. L., Schubert, G., & Zhang, M. B. (2023). *Generative AI and firm values* (NBER Working Paper No. 31222). National Bureau of Economic Research. <https://doi.org/10.3386/w31222>
- Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2023). *GPTs are GPTs: An early look at the labor market impact potential of large language models* (arXiv preprint No. 2303.10130). <https://arxiv.org/abs/2303.10130>
- Engberg, E., Görg, H., Lodefalk, M., Javed, F., Längkvist, M., Monteiro, N., Kyvik Nordås, H., Pulito, G., Schroeder, S., & Tang, A. (2024). *AI unboxed and jobs: A novel measure and firm-level evidence from three countries* (IZA Discussion Paper No. 16717). Institute of Labor Economics.
- Engberg, E., Hellsten, M., Javed, F., Lodefalk, M., Sabolová, R., Schroeder, S., & Tang, A. (2025). Artificial intelligence, hiring and employment: Job postings evidence from Sweden. *Applied Economics Letters*. Advance online publication. <https://doi.org/10.1080/13504851.2025.2481673>
- Engberg, E., Koch, M., Lodefalk, M., & Schroeder, S. (2023). *Artificial intelligence, tasks, skills and wages: Worker-level evidence from Germany* (Working Paper 12/2023). Örebro University School of Business.
- Engberg, E., Lodefalk, M., Löthman, L., & Koch, M. (2026). *Same storm, different boats: Generative AI and the age gradient in hiring* (Working Paper 2/2026). Örebro University School of Business.
- Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217. <https://doi.org/10.1002/smj.3289>
- Hosseini, S. M., & Lichtinger, G. (2026). *Generative AI as seniority-biased technological change: Evidence from U.S. résumé and job posting data* (SSRN Working Paper No. 5425555). Social Science Research Network. <https://doi.org/10.2139/ssrn.5425555>
- Hu, K. (2023, February 1). ChatGPT sets record for fastest-growing user base. *Reuters*. <https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01/>

- Humlum, A., & Vestergaard, E. (2026). *Still waters, rapid currents: Early labor market transformation under generative AI* (NBER Working Paper No. 33777). National Bureau of Economic Research. <https://doi.org/10.3386/w33777>
- Jiang, W., Park, J., Xiao, R. J., & Zhang, S. (2025). *AI and the extended workday: Productivity, contracting efficiency, and distribution of rents* (NBER Working Paper No. 33536). National Bureau of Economic Research. <https://doi.org/10.3386/w33536>
- Kauhanen, A., & Rouvinen, P. (2026). *AI has not impacted the youth labor market in Finland* (ETLA Working Paper No. 135). The Research Institute of the Finnish Economy. <https://ideas.repec.org/p/rif/wpaper/135.html>
- Murray, C., Strauss, D., Burn-Murdoch, J., & Lim, S. (2025, July 24). Is AI killing graduate jobs? *Financial Times*. <https://www.ft.com/content/99b6acb7-a079-4f57-a7bd-8317c1fbb728>
- Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science*, 381(6654), 187–192. <https://doi.org/10.1126/science.adh2586>
- SACO. (2025). *Över 100 000 arbetslösa akademiker – en rapport om lågkonjunkturens effekter*. SACO.
- Sveriges Riksbank. (n.d.). *Sök räntor och valutakurser*. Retrieved May 25, 2026, from <https://www.riksbank.se/sv/statistik/rantor-och-valutakurser/sok-rantor-och-valutakurser/>
- Swedish Public Employment Service. (2026). *JobAd Links API documentation*. JobTech Development. Retrieved February 28, 2026, from <https://jobsearch.api.jobtechdev.se/>
- Klein Teeselink, B. (2025). *Generative AI and labor market outcomes: Evidence from the United Kingdom* (SSRN Working Paper No. 5516798). Social Science Research Network. <https://doi.org/10.2139/ssrn.5516798>
- Webb, M. (2020). *The impact of artificial intelligence on the labor market* (SSRN Working Paper No. 3482150). Social Science Research Network. <https://doi.org/10.2139/ssrn.3482150>
- Westby, S., & Modestino, A. S. (2025). *The impact of generative AI on job opportunities for junior software developers* (Working Paper). Northeastern University. [https://aliciasassermodestino.com/wp-content/uploads/2025/06/Impact\\_of\\_GenAI\\_on\\_SWEs\\_061625.pdf](https://aliciasassermodestino.com/wp-content/uploads/2025/06/Impact_of_GenAI_on_SWEs_061625.pdf)

# 10 Appendix

## A.1 AI disclosure statement

During the construction of this thesis, AI-based tools have been used for a few specific purposes. Firstly, AI research tools like Scispace have been used to accumulate an initial collection of relevant literature connected to the research topic. At a later stage, LLMs were used during data cleaning and programming, helping to create the pipeline from raw job posting files to a complete dataset. At the final stage of the thesis work, AI was used to check the consistency of writing in for example citations and bibliography but also editorially in finding typos or grammar problems. The most important aspect AI tools contributed to was the pipeline of data cleaning. Transforming the raw, daily, files into a complete dataset spanning several years, removing duplicates, reformatting JSON etc. using python was beyond our capabilities as of beginning this research project. In light of the discussions on hallucinated references we have not used general LLM tools such as ChatGPT or Claude to provide sources for any empirical findings, instead opting for tools such as Scispace. Although we both have experienced this insight in other aspects of school and work, we realize the productivity increasing abilities of the technology, allowing us to do things we otherwise couldn't have done programming-wise.

## A.2 Data – Job postings

*Table A1: Most common occupational groups*

| # | SSYK lvl 1 | N         | %     | Description                                                      |
|---|------------|-----------|-------|------------------------------------------------------------------|
| 1 | 2***       | 1,345,925 | 37.9% | Occupations requiring advanced higher education qualifications   |
| 2 | 5***       | 647,792   | 18.2% | Service, care, and sales occupations                             |
| 3 | 3***       | 563,177   | 15.9% | Occupations requiring higher education or equivalent             |
| 4 | 4***       | 372,067   | 10.5% | Occupations in administration or customer service                |
| 5 | 7***       | 208,688   | 5.9%  | Occupations in construction or manufacturing                     |
| 6 | 9***       | 186,728   | 5.3%  | Occupations requiring shorter education or introductory training |
| 7 | 8***       | 118,250   | 3.3%  | Occupations in machine-based manufacturing and transport, etc.   |
| 8 | 1***       | 99,406    | 2.8%  | Managerial occupations                                           |
| 9 | 6***       | 8,918     | 0.3%  | Occupations in agriculture, horticulture, forestry, and fishing  |

*Notes:* This table presents the most common occupational groups in the full dataset. Data source: JobTech Links (2026)

**Table A2: Average AI-exposure by occupational groups**

| SSYK lvl 1 | Mean DAIOE | SD DAIOE | N       |
|------------|------------|----------|---------|
| 1          | 0.666      | 0.0318   | 69,200  |
| 2          | 0.891      | 0.161    | 205,080 |
| 3          | 0.690      | 0.128    | 100,271 |
| 4          | 0.776      | 0.167    | 49,821  |
| 5          | 0.370      | 0.079    | 130,908 |
| 6          | 0.176      | 0.085    | 656     |
| 7          | 0.200      | 0.106    | 13,974  |
| 8          | 0.220      | 0.064    | 8,174   |
| 9          | 0.070      | 0.070    | 15,922  |

Notes: This table presents DAIOE AI exposure for each occupational group. Data source: JobTech Links (2026)

### A.3 Empirical strategy – Seniority keywords

**Table A3: Number of senior job postings by keyword**

| Keyword         | N       | Keyword          | N     |
|-----------------|---------|------------------|-------|
| Chef            | 132,352 | Teamchef         | 2,222 |
| Manager         | 65,907  | Supervisor       | 1,663 |
| Senior          | 50,772  | CFO              | 1,499 |
| Specialist      | 38,709  | Regionchef       | 1,177 |
| Erfaren         | 30,865  | Chief            | 1,171 |
| Lead            | 13,490  | Years_experience | 1,097 |
| Med erfarenhet  | 9,700   | Principal        | 868   |
| Partner         | 8,599   | Sr               | 733   |
| Erfarna         | 8,326   | CTO              | 586   |
| Verksamhetschef | 6,800   | VP               | 509   |
| Head            | 6,516   | Direktör         | 315   |
| Experienced     | 6,077   | President        | 288   |
| Avdelningschef  | 5,910   | Huvud            | 252   |
| Ansvarig        | 5,420   | COO              | 223   |
| Expert          | 4,412   | CEO              | 180   |
| Ledare          | 4,326   | Skilled          | 157   |
| Director        | 4,210   | Expertise        | 134   |
| Executive       | 3,611   | Managing         | 115   |
| Rådgivare       | 2,911   | CMO              | 105   |
| Kvalificerad    | 2,636   | SVP/EVP          | 81    |
| VD              | 2,575   | Qualified        | 79    |

Notes: This table presents the number of job postings flagged for seniority by keyword. Data source: JobTech Links (2026)



**Table A4: Number of junior job postings by keyword**

| Keyword         | N      | Keyword        | N     |
|-----------------|--------|----------------|-------|
| Assistent       | 84,048 | Nyexaminerad   | 1,527 |
| Sommarjobb      | 37,065 | Lärling        | 828   |
| Junior          | 33,178 | Praktikant     | 754   |
| Student         | 19,999 | Jr             | 87    |
| Assistant       | 9,191  | Examensarbete  | 61    |
| Utan erfarenhet | 5,994  | Nyutexaminerad | 3     |
| Intern          | 4,781  | Entry level    | 2     |
| Praktik         | 3,077  | Apprentice     | 3     |
| Trainee         | 2,835  |                |       |

*Notes:* This table presents the number of job postings flagged for seniority by keyword. Data source: JobTech Links (2026)

**Table A5: Top-10 occupations in full dataset and by seniority**

| #  | Full Dataset |        |       | Junior Postings |        |       | Senior Postings |        |       |
|----|--------------|--------|-------|-----------------|--------|-------|-----------------|--------|-------|
|    | SSYK         | N      | %     | SSYK            | N      | %     | SSYK            | N      | %     |
| 1  | 2512         | 88,859 | 15.0% | 5343            | 81,363 | 40.7% | 2512            | 73,714 | 18.7% |
| 2  | 5343         | 84,143 | 14.2% | 2512            | 15,145 | 7.6%  | 3322            | 34,127 | 8.7%  |
| 3  | 3322         | 44,694 | 7.5%  | 3322            | 10,567 | 5.3%  | 2423            | 16,859 | 4.3%  |
| 4  | 2423         | 19,290 | 3.2%  | 4119            | 6,166  | 3.1%  | 2431            | 11,857 | 3.0%  |
| 5  | 4222         | 14,874 | 2.5%  | 4322            | 6,091  | 3.0%  | 1510            | 11,409 | 2.9%  |
| 6  | 2341         | 14,096 | 2.4%  | 4222            | 5,941  | 3.0%  | 3313            | 9,136  | 2.3%  |
| 7  | 1510         | 11,448 | 1.9%  | 5241            | 4,314  | 2.2%  | 4222            | 8,933  | 2.3%  |
| 8  | 4119         | 11,337 | 1.9%  | 5222            | 3,938  | 2.0%  | 1730            | 8,705  | 2.2%  |
| 9  | 3313         | 11,126 | 1.9%  | 9412            | 3,819  | 1.9%  | 1230            | 8,354  | 2.1%  |
| 10 | 4322         | 10,341 | 1.7%  | 5223            | 2,728  | 1.4%  | 2221            | 7,186  | 1.8%  |

*Notes:* This table presents the most common occupations for the full dataset, junior dataset, and senior dataset. Data source: JobTech Links (2026)

## A.4 Theoretical Framework – Derivation of the comparative Static

This is the derivation of the comparative static presented in Section 3.2, showing how an increase in AI productivity  $\phi$  affects the optimal ratio of junior to senior labor demand within a task.

### Task-level production technology

Within each task  $i$ , output is produced using a CES technology combining junior labor  $J(i)$ , senior labor  $S(i)$  and AI  $A(i)$ :

$$T(i) = [a_J(i)J(i)^\rho + a_S(i)S(i)^\rho]^{1/\rho}$$

Where  $\eta = \frac{1}{1-\rho} > 1$  is the elasticity of substitution between inputs and  $a_k(i)$ , the productivity of input  $k$  in task  $i$ .

### Augmented productivity

From the AI productivity shock  $\phi$ , the productivity of each input is:

$$a_A(i) = \phi \times r(i)$$

$$a_J(i) = \bar{a}_J(i) + \gamma_J \times a_A(i)$$

$$a_S(i) = \bar{a}_S(i) + \gamma_S \times a_A(i)$$

Where  $\bar{a}_J(i)$  and  $\bar{a}_S(i)$  is the baseline productivity before the shock.  $\gamma_J, \gamma_S > 0$  captures the degree of junior and senior workers benefit from AI augmentation.  $r(i)$  is the level of routine in a task. The publish release of ChatGPT corresponds to the increase in  $\phi$ .

### Cost minimization

The firm minimizes total costs:

$$\min_{J(i), S(i)} \int_0^1 [\omega_J J(i) + \omega_S S(i)] di$$

with regards to the production function above. Taking the ratio of the first order conditions for  $J(i)$  and  $S(i)$ :

$$\frac{w_S}{w_J} = \frac{a_J(i)}{a_S(i)} \left( \frac{J(i)}{S(i)} \right)^{\rho-1}$$

Solving for the optimal junior-to-senior ratio:

$$\frac{J(i)}{S(i)} = \left( \frac{a_J(i)}{a_S(i)} \right)^\eta \left( \frac{w_S}{w_J} \right)^\eta$$

### Log transformation

$$\ln \left[ \frac{J(i)}{S(i)} \right] = \eta \ln \left[ \frac{\bar{a}_J(i) + \gamma_J \times \phi \times r(i)}{\bar{a}_S(i) + \gamma_S \times \phi \times r(i)} \right] + \eta \ln \left[ \frac{w_S}{w_J} \right]$$

### Comparative static

Differentiating with respect to  $\phi$  and holding wages fixed:

$$\frac{\partial \ln[J(i)/S(i)]}{\partial \phi} = \eta \times r(i) \times \left[ \frac{\gamma_J}{\bar{a}_J(i) + \gamma_J \times \phi \times r(i)} - \frac{\gamma_S}{\bar{a}_S(i) + \gamma_S \times \phi \times r(i)} \right]$$