The rise of generative AI: modelling exposure, substitution, and inequality effects on the US labour market*

Raphael Auer† David Köpfer‡ Josef Švéda§

August 21, 2024

Abstract

How exposed is the labour market to ever-advancing AI capabilities, to what extent does this substitute human labour, and how will it affect inequality? We address these questions in a simulation of 711 US occupations classified by the importance and level of cognitive skills. We base our simulations on the notion that AI can only perform skills that are within its capabilities and involve computer interaction. At low AI capabilities, 7% of skills are exposed to AI uniformly across the wage spectrum. At moderate and high AI capabilities, 17% and 36% of skills are exposed on average, and up to 45% in the highest wage quartile. Examining complementary versus substitution, we model the impact on side versus core occupational skills. For example, AI capable of bookkeeping helps doctors with administrative work, freeing up time for medical examinations, but risks the jobs of bookkeepers. We find that low AI capabilities complement all workers, as side skills are simpler than core skills. However, as AI capabilities advance, core skills in lower-wage jobs become exposed, threatening substitution and increased inequality. In contrast to the intuitive notion that the rise of AI may harm white-collar workers, we find that those remain safe longer as their core skills are hard to automate.

JEL codes: E24, E51, G21, G28, J23, J24, M48, O30, O33

Keywords: Labour market, Artificial intelligence, Employment, Inequality, Automation, ChatGPT, GPT, LLM, Wage, Technology

*We thank Ryan Banerjee, Sebastian Doerr, Fiorella De Fiore, Fernando Perez-Cruz, Andras Valko, and seminar participants at the BIS for comments and suggestions. We acknowledge the use of GPT4 for editing and as mentioned in the paper. The presented views are not necessarily those of the BIS.

†Bank for International Settlements and CEPR, raphael.auer@bis.org

‡Bank for International Settlements, david.koepfer@bis.org

§Bank for International Settlements, josef.sveda@bis.org

1 Introduction

How will the advancement of generative AI complement and substitute different kinds of human labour? Recent breakthroughs have enabled generative AI to mimic human cognitive abilities in many fields, including in "white collar" professions such as law, medicine, or science. Ongoing advances and integration of the technology into day-to-day applications and workflows raise urgent policy questions.

Understanding how the potential evolution of AI will complement or substitute human skills is essential for shaping policies to ensure equitable growth and employment stability. The literature has focused on the occupation-level impact of current AI models,[1] experimental evidence of productivity impacts (Noy & Zhang, 2023; Brynjolfsson *et al.*, 2023; Peng *et al.*, 2023), and the potential for complementarity and substitution effects of AI technology at a particular state of AI development (Pizzinelli *et al.*, 2023; Acemoglu & Restrepo, 2019, 2018c,a). Except for certain types of freelancers (see e.g. Webb 2020), the broader impact of AI capabilities on the labour market yet remains to be demonstrated.

In this paper, we take a forward-looking approach: we ask the question of "what if" and examine how an AI of a hypothetical level of capabilities was to expose[2] different occupations. To shed the first light on the future impact, we build a parsimonious bottom-up quantification with a special focus on income distribution.

Our analysis proceeds in two steps. In the first step, we build on Eloundou *et al.* (2023); Felten *et al.* (2021); Gmyrek *et al.* (2023); Pizzinelli *et al.* (2023); Acemoglu (2024) and model the exposure to the technology as the capabilities of AI increase.[3] In the second step, we examine how these developments could complement or substitute human labour through the lens of their impact on core and side skills.

In the first step, we argue that the near-term impact of AI is limited a) to computer-related interactions and b) by the difficulty of the skills that AI can substitute for. In this, we only quantify the impact on skills involving interaction with a computer. We, hence, do not take into account the impact of AI on robotics that may substitute for physical work or even social interactions.[4]

Our first departure from the literature is to employ an underused part of the O*NET database that classifies skills by their difficulty. Intuitively, an AI of a certain capability level can only perform tasks up to a corresponding skill level. As the capabilities of AI advance, an

[1] See i.e.Webb (2020); Felten *et al.* (2021); Tolan *et al.* (2021); Gmyrek *et al.* (2023); Yang (2022)

[2] Throughout the paper, we use the terms "expose" and "exposure" in a neutral manner, to imply that some parts of a skill, task, or occupation could be enhanced, performed, or otherwise be affected by an AI.

[3] Similar to these approaches, we take a partial equilibrium perspective and do not take into account the interplay between skills, relative wages, human capital formation and directed technological change (Acemoglu & Restrepo, 2018c).

[4] This is in line with Acemoglu (2024), who argues that "AI is nowhere close to being able to perform most manual or social tasks", and we thus assume that it can only perform computer interactions.

increasing share of cognitive skills will hence be exposed to the technology.[5]

We nest this notion of AI capability and skill difficulty in a quantitative simulation of 711 US occupations from the O*NET database classified by the importance and the required level of cognitive skills that involve computer interactions. The model predicts that an AI capable of substituting for simple cognitive tasks – such as the minimal communication skills required for a truck driver – will expose around 7% of all skills. At low levels of AI capability, this effect holds uniformly across the entire wage spectrum, but for heterogeneous reasons. For low-income workers, a substantial share of cognitive computer skills is exposed, but the overall share of time spent on computer interaction is low. For high-income workers, only a small share of cognitive computer skills is exposed because of the larger skill requirement. However, the share of time spent using such skills is higher.[6]

As AI capabilities increase, we observe a profound difference in occupational exposure: up to 45% in the upper quartile of the wage distribution are exposed, whereas the exposure of the lowest quartile is around 26%.

What does this mean for the income distribution? We note that in line with the literature, "exposure" has a neutral meaning in that some parts of a skill, task, or job could be performed by an AI. This may lead to substitution but could also complement via increased productivity.[7]

To shed light on these issues, in the second departure from the literature and step of our simulations, we examine the extent to which AI might complement or substitute human labour. We focus on the differential impact on core versus side occupational skills, arguing that AI would tend to complement occupations wherever the auxiliary (side) skills necessary for the profession are within its capabilities. For example, if AI can organise meetings, billing, or bookkeeping for lawyers, medical doctors, or scientists, this frees up time that can be spent on core activities and thus increases productivity. However, a profession may be at risk if the core activity itself can be performed by the AI.

This exercise suggests that AI may initially complement all professions, as side skills are

[5]We take no position on how fast the evolution of the technology will materialise. Some have argued that AI may soon have dramatic impacts on the labour market (ie Korinek & Juelfs (2022)). Others argue that future advancement of AI may materialise much slower than expected. For example, Acemoglu (2024) argues that early evidence is from easy-to-learn tasks with clear outcomes (that AI can optimise for), whereas more profound productivity impacts in more subtle contexts may materialise much slower. Perez-Cruz & Shin (2024) argue that current LLMs are limited in their understanding of human interaction and higher-order beliefs.

[6]For these examples, "simple cognitive tasks" correspond to those requiring a skill level of 2.0 in the O*NET database, for example, the minimum social perceptiveness skills required for pile drivers or the minimum speaking skills required for industrial truck operators. "Medium cognitive tasks" correspond to those requiring a skill level of 3.0, for example, problem-solving skills of medical appliance technicians or the operations monitoring skills of registered nurses. "High cognitive tasks" correspond to those requiring a skill level of 4.0, for example, the persuasion skills of psychiatrists or the active listening skills of air traffic controllers.

[7]Svanberg *et al.* (2024) further note that "exposure" does not mean automation: they survey workers with "end-use" tasks to get a sense of the requirements for automation, and second, they model the cost of a model capable of meeting the requirements. Focusing on the automatability of vision, find that only 23% of occupations that are "exposed" in the sense of Eloundou *et al.* (2023); Felten *et al.* (2021) could today be automated economically. We note that our measure of exposure is more nuanced than the one in Eloundou *et al.* (2023); Felten *et al.* (2021) as we restrict the impact to skills involving computer interaction and not only model whether a skill in principle could be automated but also whether the capability level of the AI is sufficient for such automation.

generally less difficult than core skills. For example, an AI only capable of performing simple cognitive tasks has negligible exposure to core skills, whereas it, on average, exposes around 12% of side skills. However, already for moderate AI capabilities, there is divergence across the wage spectrum, with the core cognitive skills of the low-wage workers becoming roughly as exposed to AI as their side skills. In contrast, the upper quartile of the wage distribution still sees negligible exposure of core skills (5%), whereas side skills are exposed substantially (27%).

If AI capabilities are high, around 25% of both side and core skills of the lowest quartile of the wage distribution are exposed. In contrast, only 20% of the core but a staggering 62% of the side skills of the highest quartile of the income distribution become exposed.

On balance, our modelling of the impact on side and core skills hence reverses the notion that generative AI might decrease inequality in the labour market (Noy & Zhang, 2023; Brynjolfsson *et al.*, 2023). Despite being a technology that is exposing white-collar jobs more intensively, this effect is focused on the side skills of their professions, while the core skills are not in reach.[8] In contrast, a capable AI will also expose the core skills of lower-income workers, thus threatening substitution and widening inequality.

The balance of this paper is as follows: we relate our approach to the literature in Section Section 2. Next, Section 3 presents the methodology describing the evolutionary impact of ever-improving AI on occupations. It also serves as an AI exposure dependent on AI's capabilities. Thereafter, we split the AI exposure based on core and side skills Section 4 that are then used to identify complementarity and substitutional effects for individual occupations. Section 5 presents additional robustness analysis, while Section 6 concludes.

2 Literature review

Historically, technological advancements have been met with both optimism and concern regarding their implications for the labour market (Bessen, 2016). The advent of AI and machine learning technologies, in general, has intensified these debates, with researchers seeking to understand how these new tools can reshape the labour market and how the impact can differ from previous technological advancements in robotisation or computerisation (Autor, 2015).

Several recent studies have directly addressed the potential of the latest advancements in AI to significantly impact the current structure of the labour market. Brynjolfsson *et al.* (2018) argue that most occupations in the US include at least some tasks that are suitable for machine learning applications, and Eloundou *et al.* (2023) suggests that 80% of the workforce could be affected by Generative Predictive Transformers (GPTs). While these estimates are staggering, Arntz *et al.* (2016) argue that the actual vulnerability of jobs to automation is lower when considering the nuanced skills within occupations. Nonetheless, the proliferation of the latest LLMs seems to be non-negligent; Eloundou *et al.* (2023) further find 19% of US workers in the

[8] Of course, once the capability of the AI becomes extremely high such that all skills are within reach, this effect abates, and all cognitive workers are in danger of replacement.

US may see at least half of their skills impacted and Hatzius *et al.* (2023) finds 25% of current work skills in US automatable.

Current AI capabilities, in some instances, fall short of profound reasoning skills (Perez-Cruz & Shin, 2024). However, an important issue regards how the *future evolution of AI capabilities* can enhance labour productivity or crowd out workers. Recent experiments with the latest generation of AI show that it can have a positive effect in specific occupations while reducing differences among workers with varying experience levels. Noy & Zhang (2023) demonstrated that the use of ChatGPT significantly increases average productivity measured by time spent on tasks and reduces differences between high- and low-skilled workers. Brynjolfsson *et al.* (2023) studied the introduction of genAI assistant to the customer support agents and found a significantly higher number of completed tasks that were more pronounced for novice and low-skilled workers. Peng *et al.* (2023) suggests coders with access to genAI are capable of completing coding-oriented tasks up to 55% faster. AI tools can also serve as the tool to discover potential improvements in business systems (Cockburn *et al.*, 2018; Cheng *et al.*, 2022).

However, an increase in labour productivity means that less human capital is needed to maintain the same output, which could lead to layoffs or wage reductions (Acemoglu & Restrepo, 2020). In this context, Frey & Osborne (2017) predicted that up to 47% of US employment is at high risk of computerisation. Arntz *et al.* (2016) however uses a different methodology and estimates an impact of only 9%.Gmyrek *et al.* (2023) find that genAI could automate 5.1% of total employment in high-income countries, whereas low-income countries are not so susceptible. The potential for augmentation is similarly distributed across countries relative to their income levels, although the potential to augment is much larger (around four to five times). Noy & Zhang (2023) claim that ChatGPT mostly substitutes for worker effort rather than purely complementing worker skills. Yang (2022) also shows that AI can positively affect productivity and employment but adversely affects the employment of less knowledgeable workers. Some studies additionally debate the effects relative to gender (Eloundou *et al.*, 2023; Webb, 2020; Gmyrek *et al.*, 2023; Aldasoro *et al.*, 2024).

Historical experience with innovation shows that in the long-term, the displacement can be offset by an increase in the range of goods and services offered, see (Autor, 2015; Acemoglu & Restrepo, 2019). For example, Bessen (2016) shows US labour demand has increased faster in computerised occupations since 1980, although the computerisation led to substitution for other occupations, shifting employment and requiring new skills. Acemoglu *et al.* (2022) find increasing demand in AI-exposed occupations in the US since 2015. Automatisation in Japan and the US generated cost savings, allowing larger output in economy (Adachi *et al.*, 2024; Dekle, 2020; Acemoglu & Restrepo, 2020) that outweighed the displacement effects of human labour. Yang (2022) finds that AI technology is positively associated with productivity and employment in Taiwan's electronics industry for the 2002–2018 period. Acemoglu & Restrepo (2019), Acemoglu & Restrepo (2018a) and Acemoglu & Restrepo (2018c) then focus directly on the dynamics of displacement and reinstatement of labour due to automation. Based on data

from the US since World War II, Acemoglu & Restrepo (2019) claim that displacement effects occur intuitively, but they are counterbalanced by the creation of new tasks in which labour has a comparative advantage. These then change the task content of production in favour of labour because of a reinstatement effect followed by a rise in the labour share and labour demand. Acemoglu & Restrepo (2019) point out that the success of reinstatement is not automatic. It rather depends on additional variables such as the supply of new skills, demographics, or labour market institutions.[9]

Although previous innovations in automatisation and computerisation, on average, brought economic growth, they still reshaped the labour market and introduced new challenges in regional labour market structures that affected labour distribution across the skill distribution of markets. Autor (2019) documents these effects using US data showing that automation (together with international trade) led to the elimination of the bulk of non-college occupations, further leading to disproportionate polarisation of urban labour markets. Acemoglu & Restrepo (2022) document that between 50% and 70% of changes in the US wage structure over the last four decades are accounted for by workers specialised in routine tasks in industries experiencing rapid automation. Acemoglu & Restrepo (2020) show industrial robot adoption in the United States was negatively correlated with employment and wages. These examples pinpoint the importance of understanding the potential effects of technological advancements to navigate a smooth transition towards a new structure of the labour market.

The question remains how much the new wave of automation with AI is comparable to previous technological advancements. Previously, automation exposed predominantly manual labour through the invention of machines and robots. The transition process to robot-driven production, therefore, affected at its first stage rather lower-skilled labour (Acemoglu & Restrepo, 2018b). Evolving AI challenges, however, cognitive tasks and skills and creates a potential to affect different occupations by either complementing or substituting them. Earlier work by Autor & Dorn (2013) suggests that low-wage occupations faced higher substitution due to computerisation. In contrast, high-wage occupations were complemented by technology. Webb (2020) then focuses on the newer innovation in AI and states it is directed at high-skilled tasks, effectively affecting the higher-wage quantiles. A similar conclusion is reached by Eloundou *et al.* (2023) and Pizzinelli *et al.* (2023). Webb (2020) argues that the impact of AI is different from the effects of software innovation, which exposed mid-wage occupations (in line with Michaels *et al.* (2014)). Pizzinelli *et al.* (2023) emphasise high complementarity in the upper tail of the earnings distribution by AI, leading to a productivity boost instead of job displacements. The effects of AI also differ geographically. Pizzinelli *et al.* (2023); Gmyrek *et al.* (2023); Albanesi *et al.* (2023) show that more developed countries are more exposed to AI as their labour markets are more oriented to cognitive tasks. However, as AI significantly progresses, research also needs to account for the evolution of technology to fully understand its potential effects. Examining

[9] In a similar vein, Aldasoro *et al.* (2024) show in a general equilibrium model that the output effects of AI may primarily arise via the indirect impact on demand and associated changes in relative prices rather than via the direct initial productivity boost from AI adoption.

the impact of developing AI through the lens of wage distribution seems to be advantageous to formulate targeted policy responses (Furman & Seamans, 2019). As the advancements in AI technology progress, their interaction might change rapidly.

3 Measuring AI exposure: data and methodology

Predicting the impact of AI on the labour market is challenging, as the integration of the technology into real-life applications is still in its infancy, and only some synthetic benchmarks on the potential quality and efficiency improvements on certain aspects of work are available (see i.e. Tolan *et al.* (2021); Peng *et al.* (2023); Noy & Zhang (2023)). Particularly, the rapidly evolving capabilities of AI are a major source of uncertainty. In the face of these uncertainties, we construct a parsimonious bottom-up model centred on an "AI capability" parameter, which allows us to simulate the effects of evolving AI. The model is built on the skill and occupation level and later aggregated to the industry or wage-quantile level.

In this section, we show how we construct the AI Share Automatability (AISA) Index that depends on the sophistication of the AI (defined as "AI capability" above). This index rests on two main assumptions:

1. In the short to medium term, automation will affect occupational activities with computer interaction as opposed to social interactions or physical labour.
2. The skills required for performing the occupations are heterogeneous in their difficulty level. For a skill to be impacted in a certain occupation, its difficulty level needs to be within the capabilities of the AI.

We utilise data from O*NET version 27.2 and the 2022 Occupational Employment and Wage Statistics (OEWS) Survey from the US Bureau of Labor Statistics. These datasets detail around 800 different occupations (of which we can use 711 after joining across the skills tables and employment statistics) across 22 industries, providing average income, employment numbers, and ratings for up to 35 cognitive skills for each occupation in terms of required skill level (1-6) and importance (1-5). Furthermore, the data includes detailed task descriptions[10] for each occupation (on average, we have 24 task descriptions for each of the 711 occupations).

In the description of our model, we will use subscripts to denote the different levels of aggregation: the lowest level s for the skill, o for the occupation and the highest aggregation levels i for the industry or w for the wage quantile. The skill level $L_{o,s}$ is distinct for a given occupation o and skill s. For instance, the occupation of Biophysicists requires a level of 4.75 in the skill *mathematics*, while the importance of this skill $I_{o,s}$ is 3.88.

[10] `https://www.onetcenter.org/dictionary/21.0/text/task_statements.html` (release number 21.0)

3.1 Only computer interaction is automatable with AI

In this paper, we only examine the impact of AI on automating tasks that require skills involving computer interaction. Jobs performed on computers are, in the short and medium run, much more likely to incorporate AI applications compared to those involving physical labour. We acknowledge that also physical labour may, in the future, be prone to automation through improved machines and robotics. However, modelling the impact of such developments is out of the scope of the analysis at hand. Similarly, we expect social interaction to require higher degrees of social acceptance before widespread automation materialises. Certainly, cost-effectiveness and improved social skills of the AI will speed up the process, yet, as for physical labour, we expect longer timescales.

We construct a measure of the share of the time spent on computer interactions based on about 19,000 detailed task descriptions available in the O*NET database. Based on the descriptions of each occupation, we instructed GPT-4 to estimate the time spent with i) computer interaction, ii) social interaction, and iii) physical labour. The exact prompt is shown in the Box A1, and one example of task description is provided to the ChatGPT-4 in Table A1. Note that computer interaction represents working on a computer that commonly does not include communication via e-meetings or other similar social interaction.

In general, ChatGPT-4 proves very high comparability with conventional human-based procedures for categorisation purposes. Eloundou *et al.* (2023) uses both approaches (human- and GPT4-based) to directly identify occupational AI exposure, finding a very high correlation between human assessments and GPT4-based self-assessments.[11] Gmyrek *et al.* (2023) follows their approach employing ChatGPT in exploring genAI effects on the labour market worldwide.

We also cross-validate the results obtained with ChatGPT-4 by comparing the fraction of time spent on computer interaction to the AIOE indicator by Felten *et al.* (2021) in the Section 5.

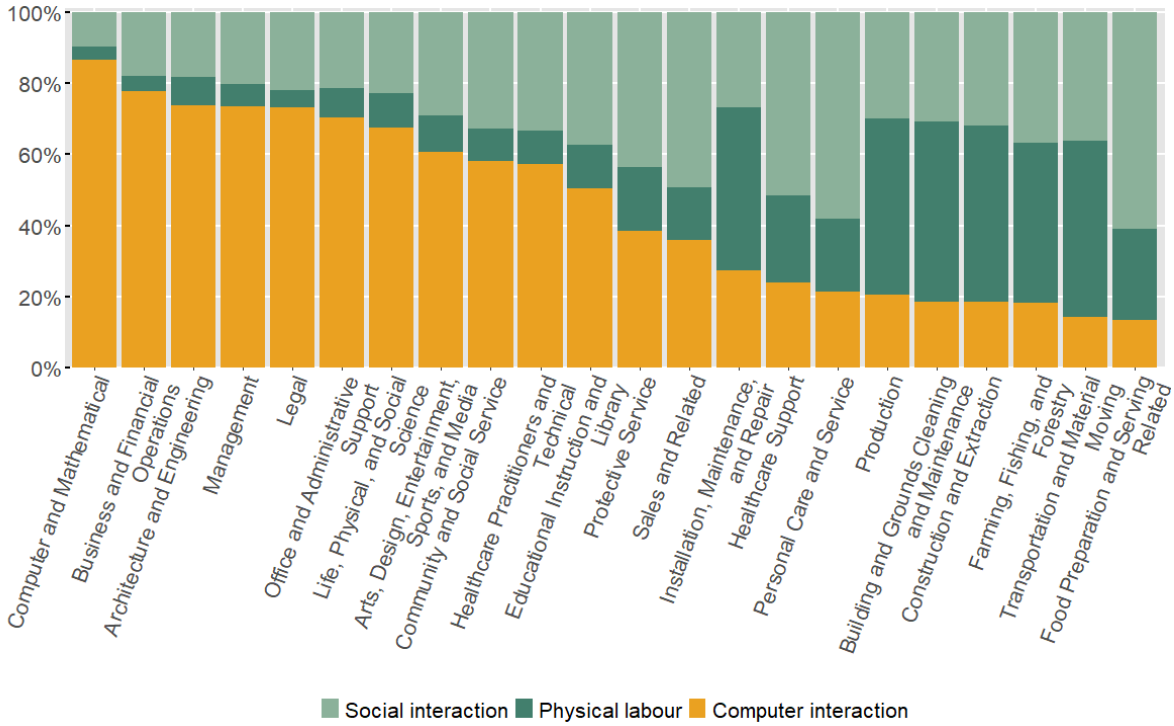
Figure 1 shows the resulting average times spent in each of these interaction types per industry (weighted by employment per occupation): Let $T_{i,o}$ denote the share of time that occupation o spends in industry i on computer interaction. The average for the industry i weighted by the employment numbers $N_{i,o}$ is calculated as:

$$T_i = \frac{\sum_{o \in O} T_{i,o} \cdot N_{i,o}}{\sum_{o \in O} N_{i,o}}$$

The distribution in Figure 1 shows that the typical office professions such as "Business and financial operations" or "Architecture and engineering" display a very large proportion in computer interactions, "Sales" and "Personal care" a large proportion of social interaction and "Production", "Construction" and "Transportation and farming" a large fraction of physical labour.

[11] Eisfeldt *et al.* (2023) further builds on their findings.

Figure 1: Time spent on technological, physical, and social interaction across industries



Note: This figure presents the fraction of time spent on i) computer interaction, ii) social interaction, and iii) physical labour (see main text for the details of the data construction) in US industries. The baseline simulations of this paper assume that only computer interaction T_i is exposed to AI.

3.2 Occupational skills need to be within the AI's capabilities to be automatable

We next measure the impact of the AI's logical capability on the exposure of various skills across occupations.[12]

The O*NET data rates the skill level and importance of around 33 cognitive skills[13] such as "reading comprehension" or "mathematics" necessary for all occupations available. Figure 2 describes their statistical properties in our dataset. The level variable "indicates the degree, or point along a continuum, to which a particular descriptor is required or needed to perform the occupation"[14] on a scale from 0 (min) to 6 (max).

The right-hand side of Figure 2 displays a histogram of (difficulty) "Level variable." This level can be interpreted as the *difficulty* level of the skill required to perform the occupation: "While the same skill can be important for a variety of occupations, the amount or level of the skill needed in those occupations can differ dramatically. For example, the skill "speaking"

[12]Note that O*NET provides several tables with data classified by level or difficulty, such as abilities and work activities tables. We test the results for those alternative tables in Section 5.

[13]The number of skills varies at each occupation between 24 and 35 with an average of 32.4 and a median of 33 skills.

[14]See O*NET website: `https://www.onetonline.org/help/online/scales`. Website accessed in June 2024.

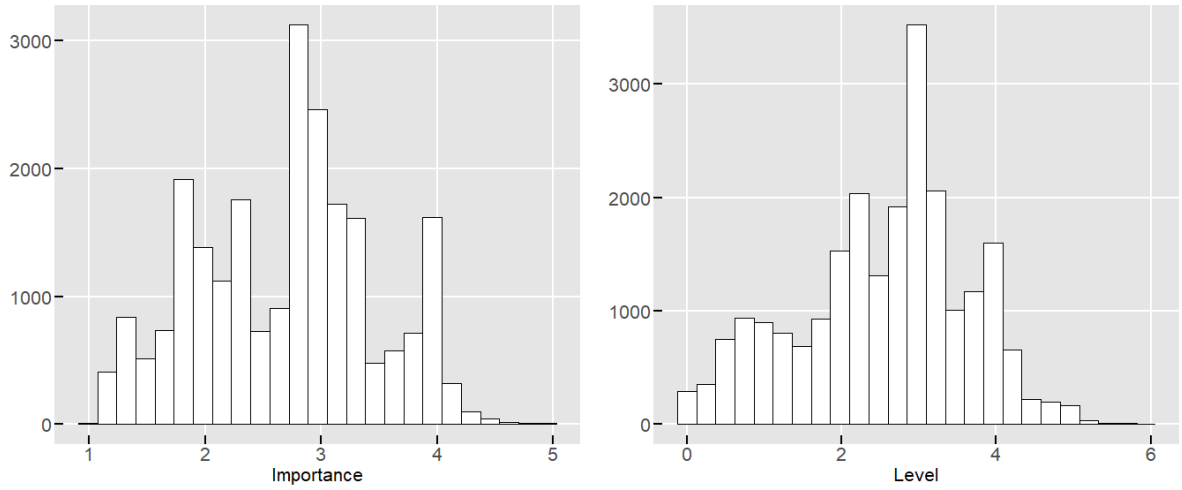
is important for both lawyers and paralegals. However, lawyers (who frequently argue cases before judges and juries) are required to have a higher level of speaking skill, while paralegals only need an average level of this skill."[15]

Table A2 in the Appendix presents the occupations most closely matching select skill levels for four cognitive skills. For example, "Troubleshooting" skills are required with a low level for "Social and Community Service Managers" whereas "Aircraft Mechanics and Service Technicians" require much higher levels to fulfil the high standards in their occupation. "Writers and Authors" do not require high skills in "Mathematics" while "Mathematicians" use it as the source of their living.

Within each occupation, skills are also classified by the "importance". This variable, "indicates the degree of importance a particular descriptor is to the occupation". The left-hand panel in Figure 2 displays a histogram of the Importance variable.[16] The possible ratings range from 1 ("Not Important") to 5 ("Extremely Important").

Figure 2: Distributions of the skill's importance and level variables

(A) Distribution of the importance $I_{i,o,s}$ **(B) Distribution of the levels $L_{i,o,s}$**



Note: The figure shows the distribution of the importance and level variables in the skills dataset from O*NET. Importance: N=23,111; Mean =2.691; SD=0.765; Min=1.04; Max=5. Level: N=23,111; Mean =2.510; SD=1.101; Min=0.060; Max=6.

In our simulations, we assume that cognitive skills that involve computer interactions are exposed once their level comes within the AI's capabilities.[17]

In the main analysis, we do not attempt to directly quantify current AI capabilities (however, we compare our results to the estimates available in the literature (subsection 3.3) and add

[15] See O*NET website: `https://www.onetonline.org/help/online/scales`. Website accessed in June 2024.

[16] See O*NET website: `https://www.onetonline.org/help/online/scales`. Website accessed in June 2024.

[17] Certainly, this is a simplification and does not take into account inherent differences in skill levels (for instance, an AI might excel at maths much sooner than on negotiation) nor the difference in time and effort to automate different skills with an AI. We also do not account for costs to the implementation that can significantly affect the decision to automate the processes, as argued in (Svanberg *et al.*, 2024).

multiple back-of-the-envelope calculations in Section 5). Rather, we introduce this as a model variable (κ_{AI}), based on which we can construct different scenarios with lower or higher AI capabilities. Although we can not foresee the speed of AI technology advancement and adoption, we are safe to assume that this parameter will grow in the future.

On a per-skill level, we construct the binary variable $A_{i,o,s}(\kappa_{AI})$ to indicate if the skill s of the occupation o in industry i is with the AI's capabilities (κ_{AI}) and thus prone to automation.

$$A_{i,o,s}(\kappa_{AI}) = \begin{cases} 0 & \text{if } L_{i,o,s} > \kappa_{AI}, \\ 1 & \text{if } L_{i,o,s} \leq \kappa_{AI}. \end{cases}$$

The overall share of skills that can be automated for a given occupation (o) as a function of the AI capability $A_{i,o}(\kappa_{AI})$ is equal to the weighted average of the impact of each of the occupation's skills, weighted by its importance ($I_{i,o,s}$) for the given occupation:

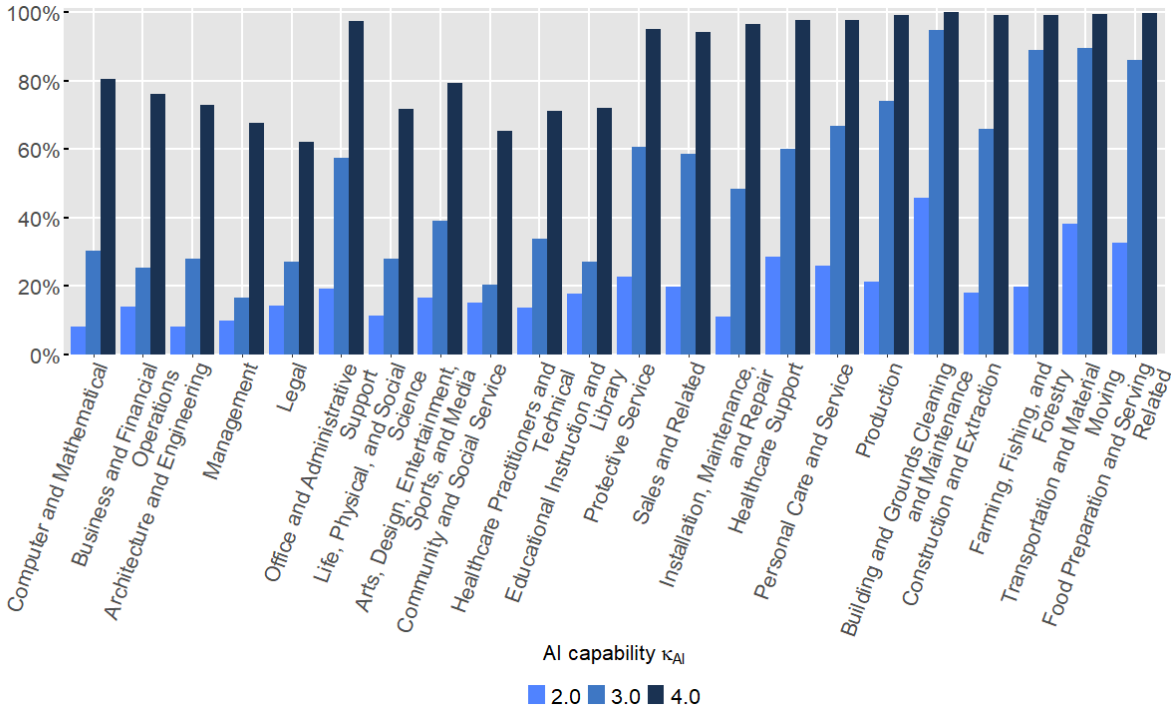
$$A_{i,o}(\kappa_{AI}) = \frac{\sum_{s \in S} A_{i,o,s}(\kappa_{AI}) \cdot I_{i,o,s}}{\sum_{s \in S} I_{i,o,s}}$$

The same share on the industry level is again calculated as an average of the shares of the individual occupations ($A_{i,o}$) weighted by the employment numbers ($N_{i,o}$) for each occupation:

$$A_{i}(\kappa_{AI}) = \frac{\sum_{o \in O} A_{i,o}(\kappa_{AI}) \cdot N_{i,o}}{\sum_{o \in O} N_{i,o}}$$

Figure 3 illustrates the share of skills within AI capabilities across industries for three low (2.0), medium (3.0) and high (4.0) AI capabilities. We see that even moderate levels of AI surpass more than 50% of the required skills necessary for many labour-intensive occupations, most strikingly, "Buildings and grounds cleaning and maintenance", where even low AI capabilities already surpass 40% of the required skills. On the opposite end, traditionally office-prone industries such as "Engineering", "Management", or "Legal" show a comparatively low share of around 20% of affected skill at medium AI levels and even at high AI capability levels does not surpass 80%.

Figure 3: Share of automatable cognitive skills given AI capabilities $A_i(\kappa_{AI})$



Note: The figure shows the share of skills $A_i(\kappa_{AI})$ within reach of the low (2.0), medium (3.0) and high (4.0) AI capability κ_{AI} in individual industries i.

3.3 AISA Index - combining computer interaction and automatability of skills

The combination of the time spent with computer interaction $T_{i,o}$ and the occupational skills within AI capabilities $A_{i,o}(\kappa_{AI})$ constitutes the AI share of automation (AISA) index:

$$AISA_{i,o}(\kappa_{AI}) = T_{i,o} \cdot A_{i,o}(\kappa_{AI})$$

Again, we aggregate the AISA index to the industry level, weighted by the employment numbers.[18]

$$AISA_i(\kappa_{AI}) = \frac{\sum_{o \in O} AISA_{i,o}(\kappa_{AI}) \cdot N_{i,o}}{\sum_{o \in O} N_{i,o}}$$

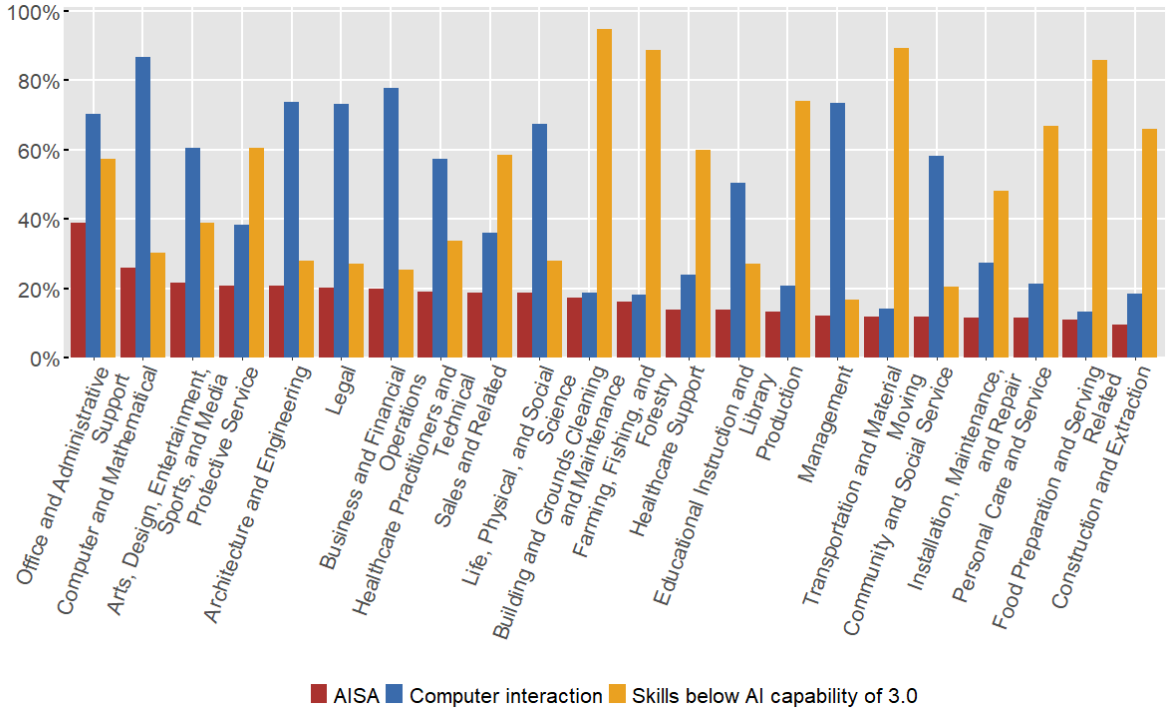
Figure 4 presents $AISA_i(\kappa_{AI} = 3.0)$ (red) within individual industries as the combination of the fraction of time spent in computer interaction T_i (yellow) and the fraction of skills within the AI's capabilities $A_i(\kappa_{AI} = 3.0)$ (blue). We see that industries like Production or Farming are "technology-limited" (i.e., despite simpler skills, much work does not involve direct computer

[18] We implicitly assume (statistic) independence between the time spent on computer interaction and individual skills. We acknowledge that this is a simplification, as some of the skills, such as "writing", are more likely performed through computer interactions than others (e.g. "negotiation"). Considering the large amount of uncertainty, we opted for this simple model. However, as more data on real-world AI usage becomes available, weighing the individual skills separately would likely be a point of refinement.

interaction, thus reducing automation potential). Conversely, industries like Legal services or engineering are "skill-level-limited"; despite high computer interaction, the complexity of skills often surpasses the AI's capabilities. At a moderate AI capability of 3.0, we observe a moderate AI impact of 10-25%[19] across the industries, with slightly greater effects in skill-level-limited industries. A noteworthy exception is the office and administrative support industry, with an almost 40% AI exposure attributable to both high computer usage and relatively low skill difficulties.

Moreover, traditionally well-paid office jobs such as Engineering or Legal have high shares of computer interaction and thus would lend themselves to easy employment of AI, assuming that the AI has the required skill levels and quality. On the other hand, there are traditionally low-paid (physical) labour-intensive jobs such as "food serving" or "cleaning", requiring less cognitive skills, but are difficult to leverage AI in their tasks because of their working modalities, which are mostly not in direct contact with a computer.

Figure 4: AISA Index at AI capability 3.0 at the industry level



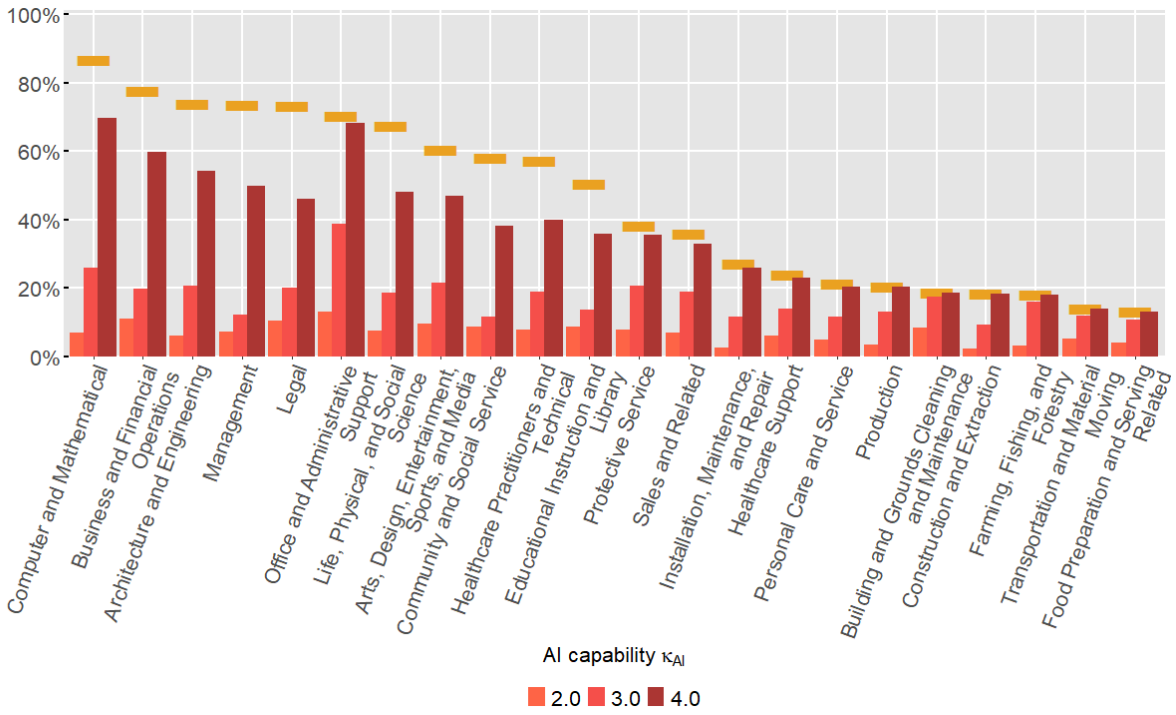
Note: The figure presents AI share automation index $AISA_i(\kappa_{AI} = 3)$ as a red bar, computer interaction T_i as a blue bar, and share of skills automatable by AI $A_i(\kappa_{AI} = 3)$ as a yellow bar across individual industries i. $AISA_{i,o}(\kappa_{AI} = 3)$ is calculated by multiplying $T_{i,o}$ and $A_{i,o}(\kappa_{AI} = 3)$ on occupational level.

Figure 5 shows the dependency of the AISA index (red bars) on the AI capability for the three AI levels (low, medium, and high). Per construction, the AISA can not surpass the fraction of time spent in computer interaction T_i (yellow lines), even at the highest AI capabilities. Industries with a high portion of computer interaction, such as Legal or Architecture,

[19]The average $AISA_{avg}(\kappa_{AI} = 3.0)$ weighted by employment numbers is 18.4%, computer interaction T_{avg} 43.3% and skills below AI capability $A_{avg}(\kappa_{AI} = 3.0)$=56.1%.

converge slowly towards their maximum, whereas occupations with low computer interaction (e.g. Production or Transport) have already almost reached their maximum at medium AI capabilities.[20]

Figure 5: $AISA_i(\kappa_{AI})$ at low (2.0), medium (3.0) and high (4.0) AI capability



Note: The figure represents the $AISA_i(\kappa_{AI})$ index for three different AI capabilities (2.0, 3.0, 4.0) at the industry level. The AISA index is dependent on the AI capability (κ_{AI}) up to a maximum level given by the computer interaction T_i (yellow lines).

Our quantification of the AISA index relates to other measures of the occupation-level exposure of AI. Felten *et al.* (2021) develop an AI-occupational exposure (AIOE) measure that gauges how important certain AI-exposed tasks are for a given occupation.[21] Pizzinelli *et al.* (2023) augment this measure (AIOE measures exposure to AI, but not whether AI complements or substitutes human labour) by a variable that measures the risk of replacement. The latter variable is constructed at the occupational level by looking into selected parts of the "work context," defined in O*NET as physical and social factors that influence the nature of work. Webb (2020), on the other hand, matches occupational task descriptions with the text of patents to match the potential impact and a similar approach was done by Yang (2022). Tolan *et al.*

[20]The overall average $AISA_{avg}(\kappa_{AI} = 2) = 7.3\%$, $AISA_{avg}(\kappa_{AI} = 3) = 18.4\%$, and $AISA_{avg}(\kappa_{AI} = 4) = 43.3\%$. The employment numbers were used as weights (see below).

[21]In more detail, Felten *et al.* (2021) combine three data sources to estimate a measure of AI occupational exposure (AIOE). First, based on information provided by the Electronic Frontier Foundation, they identify 10 applications in which AI had made "meaningful scientific progress" as of the date of writing. These applications include real-time video games, recognition and creation of speech and images, and translation. Second, these 10 applications are linked to 52 occupational abilities in the O*NET data based on a crowdsourced survey, resulting in an AI exposure measure for each of the 52 abilities (i.e., to what extent a certain ability will, in total, be exposed to the 10 applications). Lastly, AIOE is calculated as the weighted sum of the ability-level exposures, using the O*NET measures of "importance" and "prevalence" for each ability as weights.

(2021) mixes combinations of tasks and abilities with AI evaluation tasks from AI benchmarks. Brynjolfsson *et al.* (2018) chose a different approach. They define exposure by matching an established rubric with tasks and direct work activities from the O*NET database. They used a survey to establish the exposure. Eloundou *et al.* (2023) followed a similar approach but focus on a significant reduction in time to completion. They also tested an alternative approach using ChatGPT (this was also replicated by Eisfeldt *et al.* (2023)).

Our simulations raise the question regarding the current level of AI capabilities, which can be gauged from comparisons of the results of our hypothetical calculations with the survey-based ones in the literature. The red dotted line in Figure 6 displays $AISA_{avg}(\kappa_{AI})$, defined as the overall average (weighted across occupational employment as weights) as a function of AI capability (κ_{AI}):

$$AISA_{avg}(\kappa_{AI}) = \frac{\sum_{o \in O} AISA_{i,o}(\kappa_{AI}) \cdot N_{i,o}}{\sum_{o \in O} N_{i,o}}$$

For such comparisons, we note that our approach is conservative in that we assume that only computer interaction is automatable. $AISA_{avg}(\kappa_{AI})$ (red dotted line) increases gradually and tops out at 43.3%, which is the average level of computer interaction across all industries that bounds AISA.

We compare $AISA_{avg}(\kappa_{AI})$ to four datapoints found in the literature. Hatzius *et al.* (2023) and Eloundou *et al.* (2023) suggest that on average 25%[22] and 30%[23] of occupations will be exposed respectively. A level of κ_{AI} of 3.2 and 3.6, respectively, would match these numbers in the aggregate.

An earlier literature had suggested a larger automatisation potential that exceeds the maximum of $AISA_{avg}(\kappa_{AI})$ in our simulations. Frey & Osborne (2017) estimate a maximum exposure of 47%[24], while Webb (2020) estimates an exposure of 54.5%[25].

Some of these higher estimates may rely on the assumption that also types of work other than computer interaction are automatable with AI. For better comparability, Figure 6 also reports an alternative simulation for $AISA_{avg}(\kappa_{AI})$ that is discussed in the robustness Section 5. This alternative measure assumes that in addition to computer interaction, also a share of social interactions, such as communication with clients via e-meetings or taking orders in a restaurant, can be automated with AI: $T_{i,o}^{25\%Social}$ is calculated assuming that 25% of social interaction can be automated (if the skill difficulties are in reach of the AI's capabilities).

$AISA_{avg}^{T_{i,o}^{25\%Social}}(\kappa_{AI})$ is depicted by the red dashed line in Figure 6 and tops out at 51.8%.

[22]The value represents AI exposure based on an evaluation of 13 work activities of O*NET.

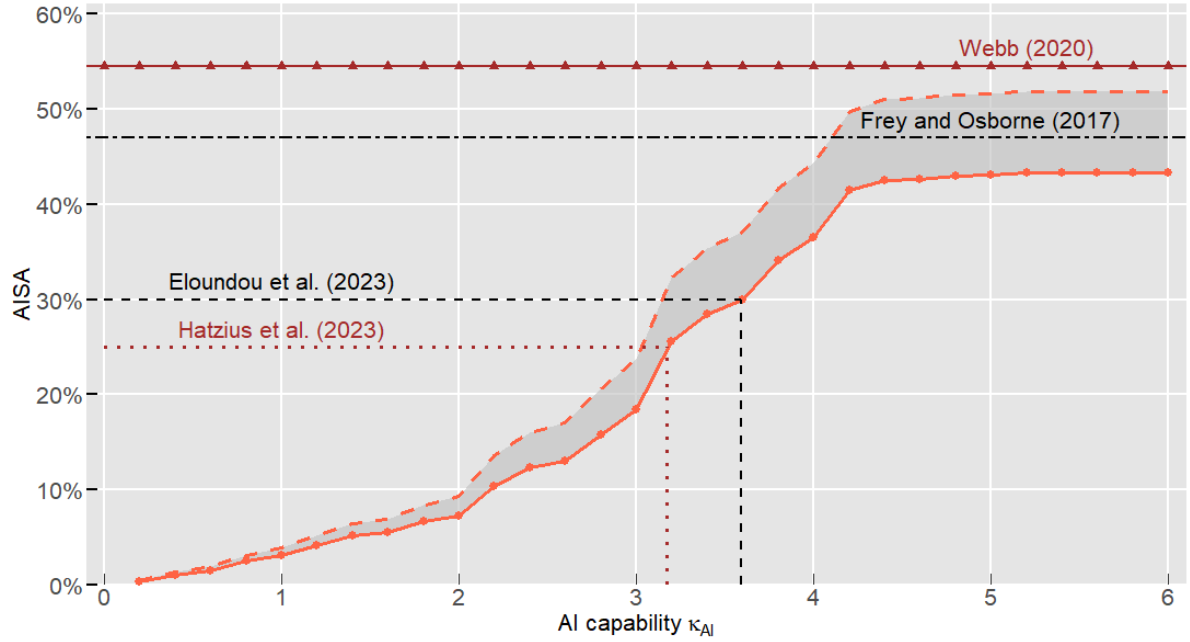
[23]This value represents the mean AI exposure based on a human assessment of the professions in which AI can reduce the time spent on tasks by at least 50% only with the use of the OpenAI's ChatGPT (alone, or in their "beta" case when it is integrated in a company's systems).

[24]Note that this figure is discussed in the literature. Arntz *et al.* (2016) revise these estimates using a modified methodology and find that only 9% of jobs in the US are at high risk of automation.

[25]The value represents the exposure to AI. Webb (2020) also models the exposure to software (50.69%) and robots (48.61%).

Using this line as a benchmark, an AI capability level of 2.9 and 3.1, respectively, would match the averages mentioned in Eloundou *et al.* (2023) and Hatzius *et al.* (2023). In contrast, reproducing the figure of Frey & Osborne (2017) would require a very high AI capability of 4.1 (while there is no level of AI capabilities that allows us to reproduce the predictions of Webb (2020)).

Figure 6: AISA as a function o f κ_{AI} and comparison to other estimates



Note:
The figure presents the level of $AISA_{avg}(\kappa_{AI})$ as a function of AI capability κ_{AI}. The average AISA is constructed as the weighted mean across industries, using employment numbers of each occupation as weights. $AISA_{avg}(\kappa_{AI})$ is represented by the red dotted line. The figure also depicts an alternative measure from the robustness exercise in subsection 5.1, assuming that AI can automate not only computer interaction but also 25% of social interactions ($T_{i,o}^{25\%Social}$). $AISA_{avg}^{T_{i,o}^{25\%Social}}(\kappa_{AI})$ is depicted by the red dashed line. The figure also adds four lines corresponding to average AI exposure levels found in the literature (Eloundou *et al.* (2023) (black dashed vertical line), Hatzius *et al.* (2023) (brown dotted vertical line), Frey & Osborne (2017) (black dot-dash horizontal line), and Webb (2020) (brown triangular horizontal line).

Figure 7 maps the AISA index across the wage spectrum. To construct the wage spectrum, we used weighted quantiles to take differences in the numbers of employees per occupation N_o into account. Assignment to a wage quantile was done by first indexing (j) all occupations in ascending order of their wage $W_{o,j}$ and calculating the cumulative sum of employees leading up to the index of the particular occupation $j = j_o$, divided by the total number of employees in the dataset:[26]

$$p_o = \frac{\sum_{j=1}^{j=j_o} N_{o,j}}{\sum_{j \in J} N_{o,j}}$$

[26]Imagine, lining up all employees based on their wage, then p_o would denote the relative position as a fraction of the entire length of the line of the last representative of a given occupation o in that line.

Based on that relative position, an occupation is assigned to wage quantile w from a total of W quantiles if its relative position is in between equally spaced boundaries:

$$\frac{w-1}{W} < p_o \leq \frac{w}{W}$$

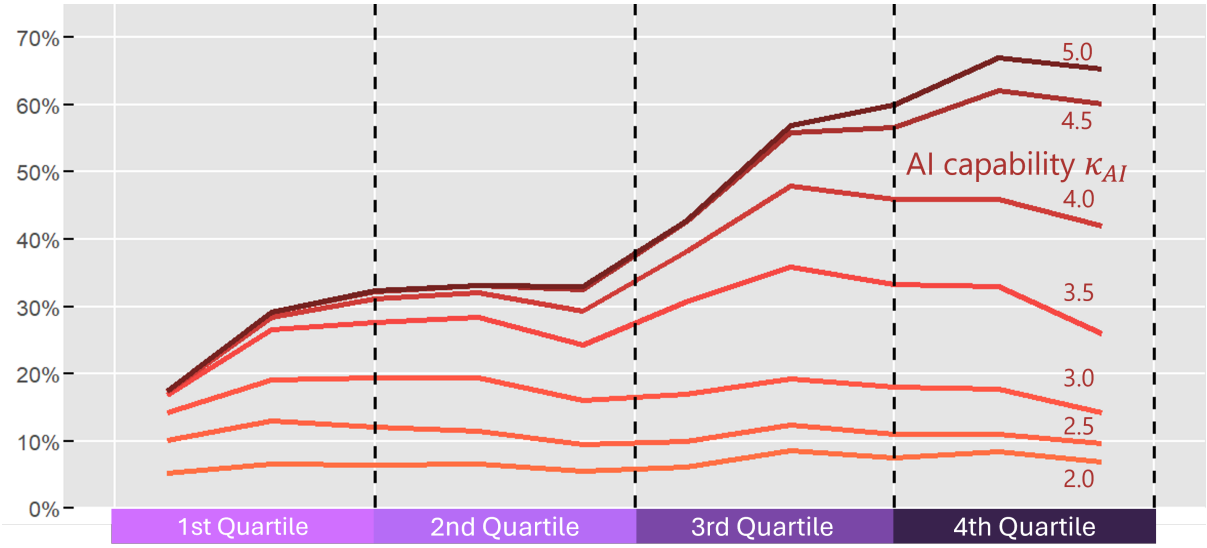
The weighted average of $AISA_w(\kappa_{AI})$ within a specific quantile w is calculated (analogously to the Industry) with an average of the per occupation $AISA_{w,o}(\kappa_{AI})$ weighted by the number of employees in that occupation $N_{w,o}$:

$$AISA_w(\kappa_{AI}) = \frac{\sum_{o \in O} AISA_{w,o}(\kappa_{AI}) \cdot N_{w,o}}{\sum_{o \in O} N_{w,o}}$$

Figure 7 demonstrates that the impact on occupations is uniform and relatively small across the wage spectrum at low levels of AI capability (below 10% at the AI capability of 2.0). However, as AI capabilities increase, particularly from AI capability 3.5 onwards, the effects begin to diverge significantly.

At the highest levels of AI capability, the additional impact is predominantly on the highest-paid occupations, as lower-paid positions have already reached their maximum potential advantage from AI integration. We note that the shape of the exposure across the wage spectrum for higher AI capabilities (i.e. around 4) mimics that fond for the survey-based estimates of Eloundou *et al.* (2023).[27]

Figure 7: AISA Index for various AI capabilities across wage quantiles



Note: The figure depicts AISA index $AISA_w(\kappa_{AI})$ for increasing AI capability κ_{AI} across wage quantiles w. In this case, $AISA_w(\kappa_{AI})$ was calculated on deciles.

[27] Eloundou *et al.* (2023) finds a similar shape where higher wage quantiles are much more exposed to AI. Their shape resembles the one depicted in Figure 7 most closely for $AISA_w(\kappa_{AI} = 4.0)$.

4 Complementing and substituting AI: the role of Core and Side skills

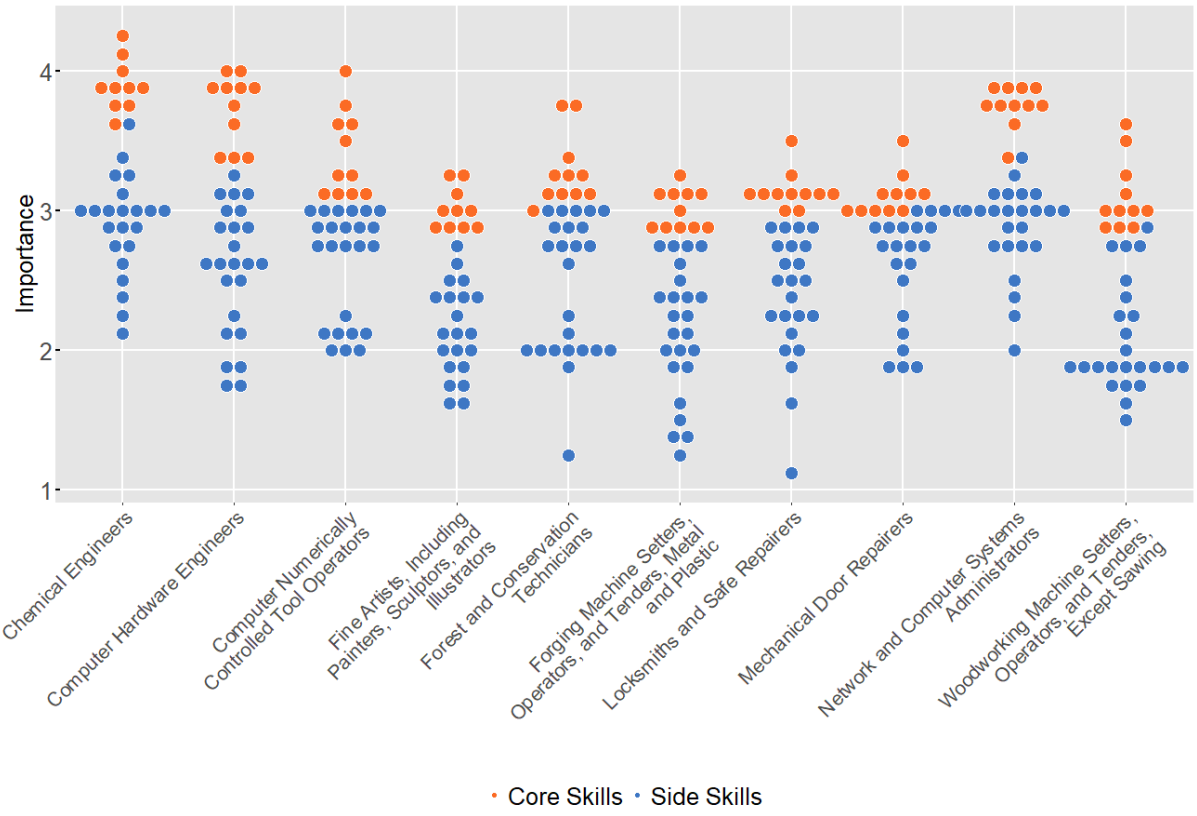
Next, we classify an occupation's skills into core and side ones and compute measures of AI exposures separately for these two groups.

4.1 Exposure of core and side cognitive skills

To model the effect of complementarity and substitution in our model, we split the cognitive skills by their importance for each individual occupation into two categories: the top third, approximately 11 skills we designate as "core skills", and the remainder, referred to as "side skills". We find support for the division from other tables available in the O*NET database that show a clear relationship between higher importance and core/side tasks[28] Figure A1). Figure 8 illustrates this division for 10 randomly selected occupations.

[28]Note that the construction of core and side tasks in O*NET database follows a very similar methodology we apply. The division is constructed based on the relevance and importance rating of the tasks where the criteria for core tasks are (a) relevance 67% and (b) a mean importance rating of 3.0.

Figure 8: Core vs. Side skills differentiation for individual occupations



Note: The figure shows the distribution of core (orange) and side (blue) skills for a random selection of occupations at the most granular level. The top third (11 out of 33 for the median number of skills per occupation) important skills are designated *core* skills, leaving the bottom 2/3 of skills as *side* skills.

As before, the AISA index is calculated as the average of the individual occupations within a wage quantile w, but this time split between core and side tasks:

$$AISA_w^{core}(\kappa_{AI}) = \frac{\sum_{o \in O} AISA_{w,o}^{core}(\kappa_{AI}) \cdot N_{w,o}}{\sum_{o \in O} N_{w,o}}$$

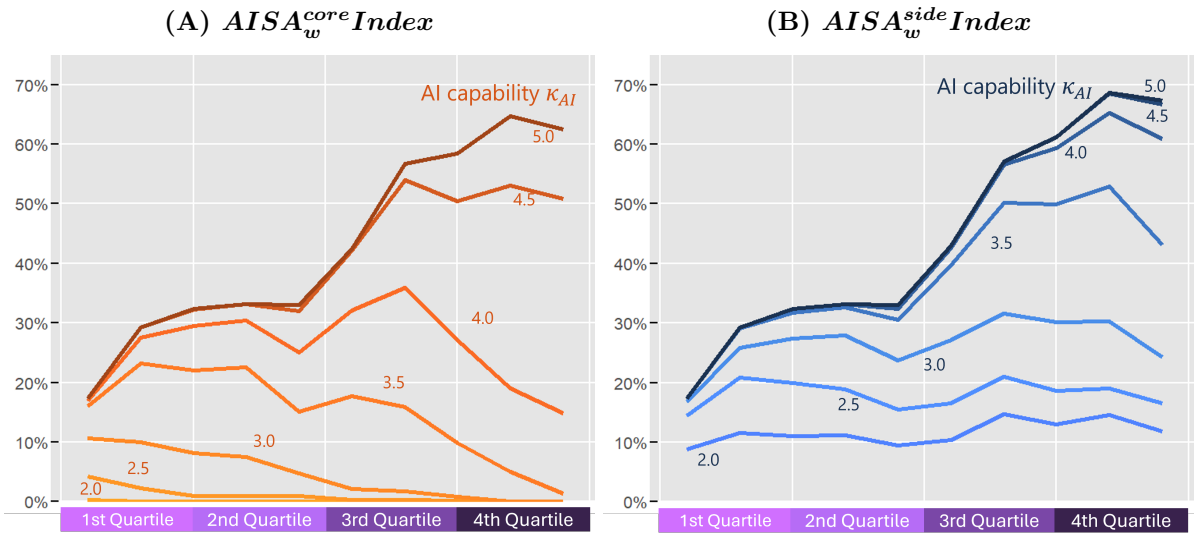
$$AISA_w^{side}(\kappa_{AI}) = \frac{\sum_{o \in O} AISA_{w,o}^{side}(\kappa_{AI}) \cdot N_{w,o}}{\sum_{o \in O} N_{w,o}}$$

Figure 9 shows AI impact for core skills (Figure 9A) and side skills (Figure 9B) separately across the wage spectrum. Comparing the two approaches, the AI impact of side skills is pronounced at much lower AI capabilities compared to core skills. This suggests that AI advancements affect side tasks more immediately across all wage quantiles. For lower-wage occupations, side skills saturate at around 20-30% total exposure, whereas core skills start with lower impact and increase more gradually, reaching a comparable level at the later-stage AI capabilities.

Importantly, occupations in higher wage brackets are influenced twice as much as those

in the mid-wage range for high AI capabilities. This disparity becomes most apparent when examining the effect on core skills. At a moderate AI capability of 3.0, the core skills (Figure 9A) in the fourth quartile interval remain largely unaffected; however, as the AI reaches a highly advanced level of 5.0, the impact on these skills escalates dramatically, soaring to 60%. These observations underscore a pivotal trend: as AI progresses, its capacity to undertake complex, high-value skills increases, disproportionately affecting higher-paid occupations.

Figure 9: Core and Side skill exposure across the wage quantiles



Note: The figure illustrates the impacts of $AISA_w^{core}(\kappa_{AI})$ (orange, Figure 9A) and $AISA_w^{side}(\kappa_{AI})$ (blue, Figure 9B) with increasing AI capability (κ_{AI}) across the wage quantiles w. In this case, we used deciles for w.

4.2 Core and side skills as a proxy for complementarity and substitution

A consistent pattern emerges upon revisiting the impact of AI across core and side skills. Notably, there is a strong correlation between the skills' Importance and Level, thus side skills are already affected at lower AI capabilities than the core skills. This opens a window where the AI is competent in mastering the side skills but not yet the core skills. If we consider the AI's proficiency in executing core skills as an indicator of potential labour *substitution*:

$$R_w(\kappa_{AI}) = AISA_w^{core}(\kappa_{AI})$$

Further, we interpret situations where AI surpasses side skills but not core skills as periods of AI *complementarity* for a given occupation.

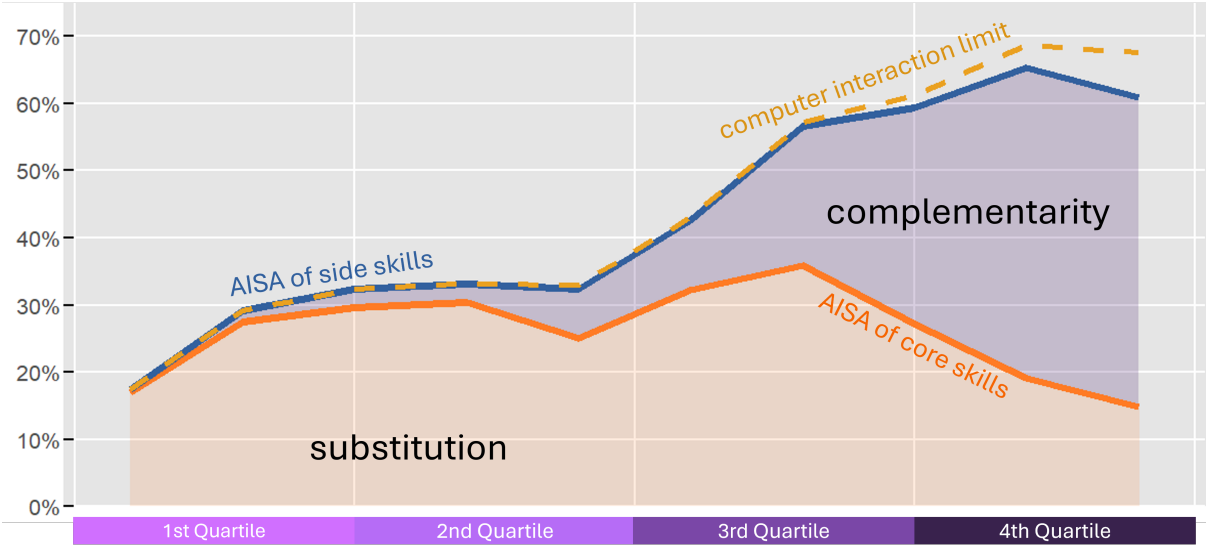
$$C_w(\kappa_{AI}) = AISA_w^{side}(\kappa_{AI}) - AISA_w^{core}(\kappa_{AI})$$

In this framework, the difference between AI's impact on core and side skills emerges as a measure of potential productivity gains for employees. This perspective allows for a nuanced

understanding of AI's role in the workplace: not merely as a substitution for human labour but as a tool for augmenting human efficiency and capability.

Figure 10 illustrates the interplay between complementarity and substitution of labour at a high AI capability of 4.0. The effects of both side and core skills are presented in one figure showing the areas where AI complements and substitutes occupations. As observed, a significant portion of higher-paying jobs (approximately 40% - after abstracting the effect of substitution) benefits from AI, indicating a substantial productivity increase. This complementarity effect is particularly evident in occupations requiring advanced side skills, where AI enhances human capabilities. Conversely, AI demonstrates a notable ability to handle core skills in lower-wage occupations. These jobs, often reliant on manual tasks, would need to be adapted to enable humans to manage new complex and creative tasks. Despite this, AI's impact on side skills in lower-wage occupations is limited to around 30%, as these roles heavily depend on manual labour.

Figure 10: Complementary and substitution percentage across the wage spectrum for an AI capability of 4.0



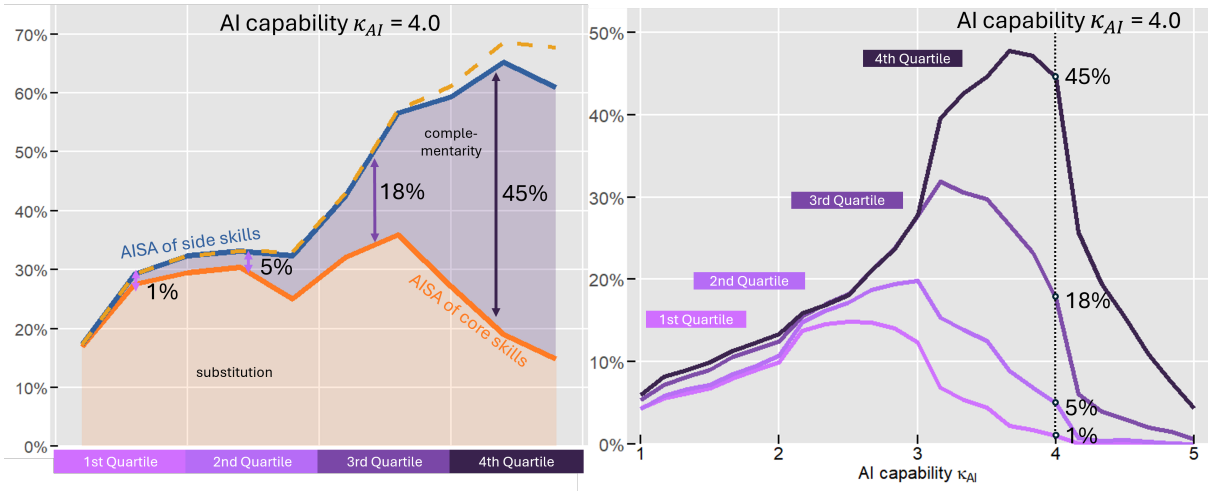
Note: The figure illustrates the amount of complementarity $C_w(\kappa_{AI})$ and substitution $R_w(\kappa_{AI})$ at a higher AI capability of 4.0 ($\kappa_{AI} = 4.0$) across the wage spectrum w. In this case, the variables are calculated on deciles. The substitution (orange area) equals $AISA_w^{core}(\kappa_{AI})$ and complementarity (purple area) equals the difference between $AISA_w^{side}(\kappa_{AI})$ and $AISA_w^{core}(\kappa_{AI})$. Orange line depicts $AISA_w^{core}(\kappa_{AI})$ and blue line $AISA_w^{side}(\kappa_{AI})$. The yellow line indicates the computer interaction limit, i.e. the maximum interaction of occupations with AI: $AISA_w(\kappa_{AI=6})$.

Switching the viewpoint to the continuing progress of the AI capabilities, Figure 11 (left-hand side) repeats Figure 10, highlighting $C_w(\kappa_{AI} = 4.0)$ as the difference between $AISA_w^{side}(\kappa_{AI})$ and $AISA_w^{core}$ for the four wage quartiles. Figure 11 (right-hand side) displays the dependence of the complementarity $C_w(\kappa_{AI})$ on the AI capabilities κ_{AI} for those wage quantiles. For each of the wage brackets, we observe distinct peaks of complementarity, starting once the AI becomes first capable of the side skills and levelling off once the AI is capable of performing the core skills.

The progression of the peaks is very interesting: For instance, the lowest wage quartile shows only a modest peak with only a peak of 15% occupations experiencing a productivity boost, which is achieved relatively "early" at lower AI capability (around 2.5). Interestingly, this peak of AI impact escalates in value and occurs at higher AI capabilities in the higher wage quantiles with the maximum complementarity effect for the highest quartile of about 50% at an AI capability as high as 3.7, after which it drops off sharply. This pattern suggests an impending period of significant efficiency gains, particularly in professional occupations, which is likely to introduce considerable disruption, especially in the high-wage labour market. The implications of this trend are profound, indicating a turbulent era where high-wage industries may experience both heightened productivity and a considerable reshaping of job roles due to advancing AI capabilities.

In contrast to Cazzaniga *et al.* (2024), we hence find that increased AI capabilities may have a rather rising effect on income inequality at the middle to high AI capabilities, as the complementarity effect will be the highest for high-wage occupations. Of course, extremely advanced AI capabilities could potentially recalibrate the labour market, especially in industries characterised by high wages and complex skills. Such high AI capabilities are associated with the automation of high-skill, high-wage jobs as their core tasks become exposed to the technology.

Figure 11: AI complementarity as AI capabilities increase



Note: The left panel presents the complementarity $C_w(\kappa_{AI})$ and substitution $R_w(\kappa_{AI})$ identification process for AI capability $\kappa_{AI} = 4.0$ from Figure 10. In addition, we include arrows that display the magnitude of the complementarity effect $C_w(\kappa_{AI} = 4.0)$ for each of the four wage quartiles w. The right panel illustrates how the complementarity effect $C_w(\kappa_{AI})$ evolves as the AI capabilities κ_{AI} change. Note that complementarity effect $C_w(\kappa_{AI})$ represents the difference between $AISA_w^{side}(\kappa_{AI})$ and $AISA_w^{core}(\kappa_{AI})$. The magnitudes of the complementarity effect are depicted as arrows in the left panel, and they also appear in the right panel when $\kappa_{AI} = 4.0$.

5 Robustness Analysis

Next, we test the robustness of our results by introducing several adjustments to our framework. First, we present several alternatives to the time spent on computer interaction $T_{i,o}$, either allowing for social interactions to be partly automatable or using a different measure from

another source. Second, we use abilities and work activities instead of skills in the O*NET database to construct AISA. Lastly, we modify the definitions of core and side skills.

5.1 Robustness with regards to the computer interaction variable

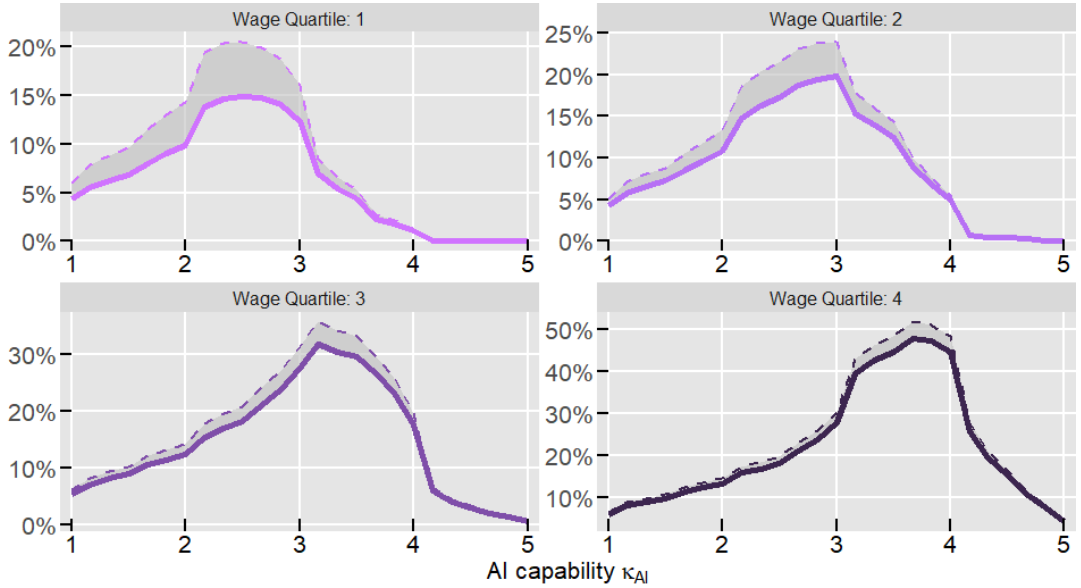
We start with robustness analysis with respect to the computer interaction variable, $T_{i,o}$. In the baseline simulations, the latter is measured rather conservatively, even when it comes to the near-term effect of AI: also, some social interactions, such as communication with clients via e-meetings or taking orders in a restaurant, might be automatable with generative AI. As a first robustness test, we hence construct a new measure $T_{i,o}^{25\%Social}$ that is equal to the time spent on computer interactions plus 25% of the time spent on social interactions:

$$T_{i,o}^{25\%Social} = T_{i,o} + 0.25 \cdot S_{i,o}$$

Here, $S_{i,o}$ is the time spent on social interaction, as constructed above in Subsection 3.1 and displayed in Figure 1.

Figure 12 compares the baseline results for our complementarity measure $C_w(\kappa_{AI})$ from Figure 11 (right panel) with the approach that considers also 25% of the time spent on social interaction automatable $T_{i,o}^{25\%Social}$. The baseline results are depicted with solid lines, and the results using $T_{i,o}^{25\%Social}$ with dashed lines. Under this alternative specification, the "hump" shape of the impact of AI is more pronounced, as the difference between $AISA_w^{side}(\kappa_{AI})$ and $AISA_w^{core}(\kappa_{AI})$ peaks at a higher level across all four wage quartiles. Additionally, this difference is more pronounced for the first and the second quartiles than for above-average wages.

Figure 12: Accounting for automation of social interactions with $T_{i,o}^{25\%Social}$



Note: The figure compares complementarity effects $C_w(\kappa_{AI})$ of the baseline model with the alternative specification that considers both computer interaction $T_{i,o}$ and 25% of social interactions $S_{i,o}$ to be automatable with AI ($T_{i,o}^{25\%Social}$). Each panel shows the evolution of $C_w(\kappa_{AI})$ for an individual wage quartile against the higher AI capabilities κ_{AI}. The solid lines represent the baseline model and are identical to the results presented in Figure 11 right panel. The dashed lines represent $C_w(\kappa_{AI})$ when $T_{i,o}^{25\%Social}$ is used instead of $T_{i,o}$.

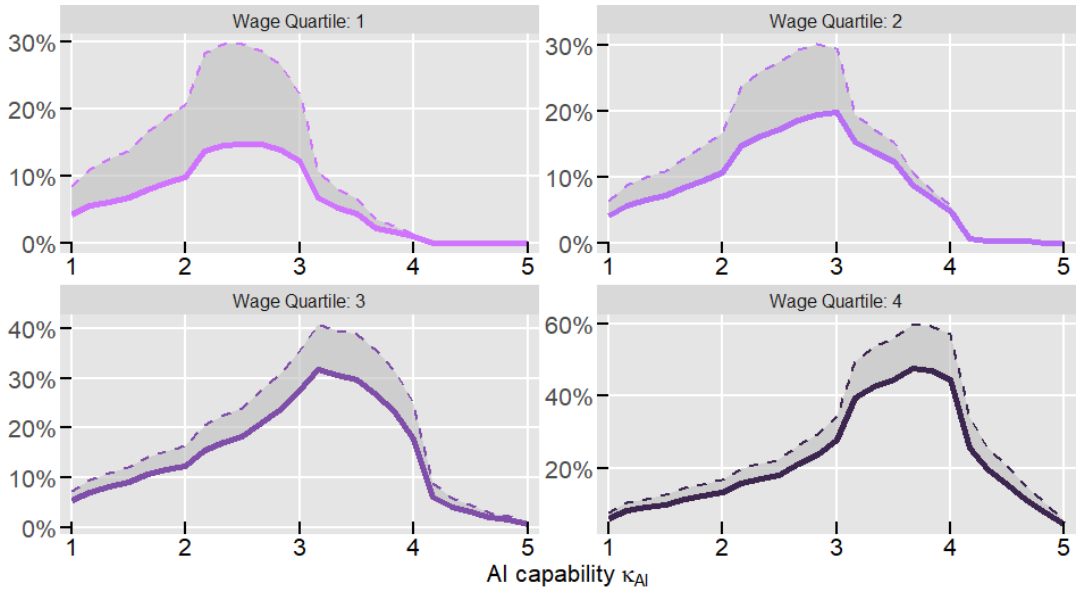
To further test the robustness of our findings, we replace the computer interaction variable with the AIOE index from Felten *et al.* (2021). The AIOE index measures occupational exposure to AI based on a survey conducted in the USA using the O*NET database, making it directly applicable in our model. Although AIOE is a normalised index focusing on the variability of AI exposure among different occupations rather than the absolute level, it is useful for understanding the relative positions of occupations in terms of AI exposure. Notably, the correlation between computer interaction and AIOE is 0.863 (see Figure A2).

To understand the absolute magnitude of the impact using AIOE, we rescaled the AIOE index to a zero-to-one interval using min-max normalisation[29]. We then re-ran our baseline estimations with this rescaled AIOE instead of $T_{i,o}$. The results are shown in Figure 13, which displays ($_o(\kappa_{AI})$ for each quartile of the occupational wage distribution.

A visual inspection reveals that the impacts are generally larger when using the rescaled AIOE. Notably, the size of the complementarity effect increases in the higher wage quartiles, similar to when using the computer interaction variable $T_{i,o}$. While the first and second quartiles remain at a maximum level of 30%, the third quartile increases to 40%, and the fourth quartile to 60%. Therefore, in line with our baseline model, AIOE also indicates that higher-wage occupations benefit more from higher AI levels.

[29]Figure A2 presents the scatterplot with the original, unscaled AIOE. In this case, AIOE falls into the negative value range, which would make the results uninterpretable.

Figure 13: Using AIOE index instead of computer interaction

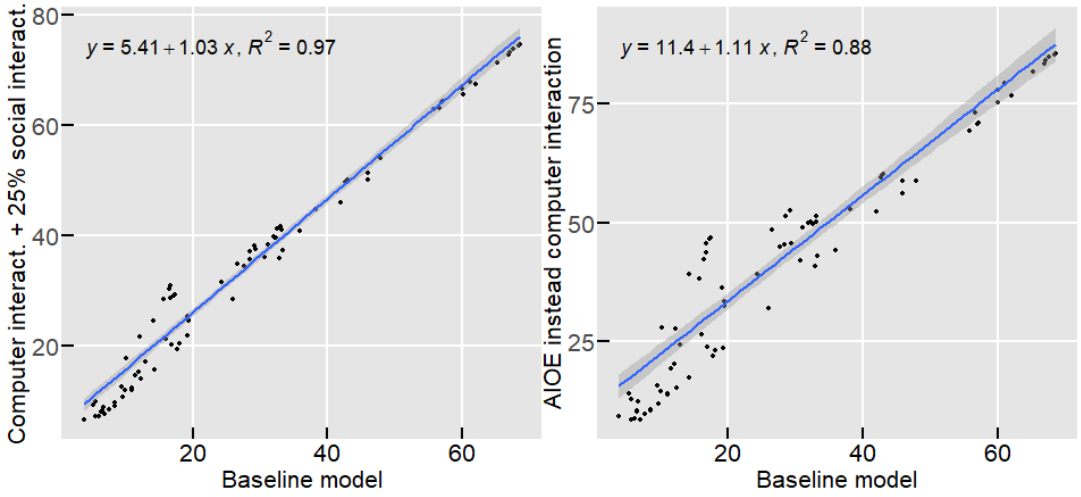


Note: The figure compares complementarity effects $C_w(\kappa_{AI})$ of the baseline model with the alternative specification that considers AIOE index from (Felten *et al.*, 2021) instead of $T_{i,o}$. AIOE was rescaled on the occupational level to zero-one interval with min-max normalisation. Each panel shows the evolution of $C_w(\kappa_{AI})$ for individual wage quartiles against the progressing AI capabilities κ_{AI}. The solid lines represent the baseline model and are identical to the results presented in Figure 11 right panel. The dashed lines represent $C_w(\kappa_{AI})$ when AIOE index is used instead of $T_{i,o}$.

Next, we compare the alternative versions of AISA to our baseline model. Figure 14 presents their scatter plots and simple regression lines. When we apply $T_{i,o}^{25\%Social}$ instead of computer interaction $T_{i,o}$, AISA is, on average, 5.41pp higher and remains stable relative to the baseline results. Using AIOE instead of computer interaction, AISA is, on average, 11pp higher[30] than the baseline model, with more pronounced magnitudes at higher values of the baseline model.

[30]Note that the absolute values of AISA with applied AIOE should not be, however, interpreted in absolute terms. The indicator only measures the variation among occupations and does not reveal the absolute impact.

Figure 14: AISA in comparison to alternative models (computer interaction oriented)



Note: The figure presents scatter plots of AISA (as used in Figure 7) from the baseline model (x-axis) and AISA calculated with the alternative setups of the baseline model used for the robustness analysis (y-axis). We consider two different setups: i) $T_{i,o}^{25\%Social}$ instead of computer interaction $T_{i,o}$, and ii) the rescaled AIOE index from (Felten *et al.*, 2021) instead of computer interaction $T_{i,o}$. Panels also display summaries of simple OLS regression lines.

5.2 Using abilities or work activities instead of skills

The second building block of our model relies on the O*NET skills dataset, which is characterised by importance $I_{o,s}$ and *difficulty* level $L_{o,s}$. Although skills offer the best interpretability for our baseline results, alternative tables with the same structure, such as those for (i) abilities and (ii) work activities, are available in the O*NET.

Abilities are defined as enduring attributes of the individual that influence performance. O*NET sorts the abilities into four categories: i) cognitive, ii) physical, iii) psychomotor, and iv) sensory abilities. For the purposes of the robustness check, and in line with the scope of our focus on the automatability of cognitive computer interactions, we only focus on the first category. Cognitive abilities are defined as "abilities that influence the acquisition and application of knowledge in problem-solving".[31] There are 21 unique cognitive abilities such as "Deductive Reasoning" or "Written Comprehension". The *difficulty* level $L_{o,s}^{abilities}$ and importance $I_{o,s}^{abilities}$ for those are defined based on the same scale as the skills table in O*NET with the level ranging from 0 to 6 and importance ranging from 1 to 5.

We use the level $L_{o,s}^{abilities}$ and importance $I_{o,s}^{abilities}$ variables to calculate an alternative indicator to shares of automatable cognitive skills for different AI capabilities $A_o(\kappa_{AI})$. $A_o^{abilities}(\kappa_{AI})$ is calculated as follows:

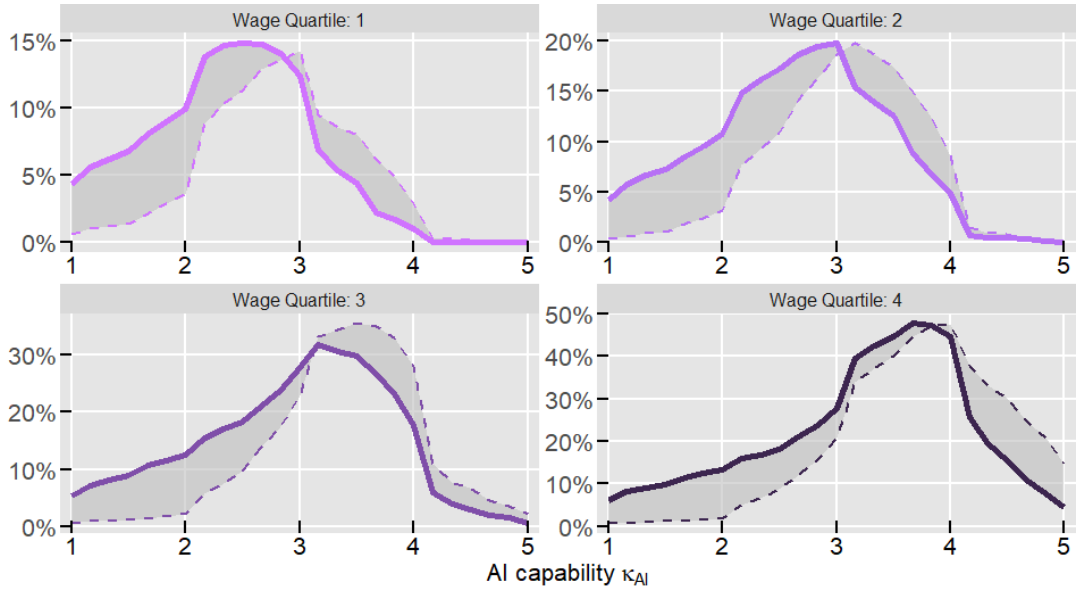
[31] see `https://www.onetonline.org/find/descriptor/browse/1.A` The website was accessed in June 2024.

$$A_o^{abilities}(\kappa_{AI}) = \frac{\sum_{s\in S} A_{o,s}^{abilities}(\kappa_{AI}) \cdot I_{o,s}^{abilities}}{\sum_{s\in S} I_{o,s}^{abilities}}; \quad A_{o,s}^{abilities}(\kappa_{AI}) = \begin{cases} 0 & \text{if } L_{o,s}^{abilities} > \kappa_{AI}, \\ 1 & \text{if } L_{o,s}^{abilities} \leq \kappa_{AI} \end{cases}$$

$A_o^{abilities}(\kappa_{AI})$ enables us to calculate a new AISA index and complementarity measure $C_o(\kappa_{AI})$ that reflects cognitive abilities instead of skills.[32]

Figure 15 compares $C_o(\kappa_{AI})$ using cognitive abilities tables instead of skills with the baseline results. Cognitive abilities (dashed lines) yield outcomes similar to the baseline model (solid lines), with peak complementarity for each quartile closely matching the baseline values and maintaining the same position relative to AI capability κ_{AI}. The average complementarity effect across all quartiles is, however, lower by 2.0pp (11.2%, compared to 13.2% in the baseline model).

Figure 15: Cognitive abilities instead of skills



Note: The figure compares complementarity effects $C_w(\kappa_{AI})$ of the baseline model with the alternative approach that uses the cognitive abilities table instead of the skills table from O*NET. Each panel shows the evolution of $C_w(\kappa_{AI})$ for individual wage quartiles against higher progressing AI capabilities κ_{AI}. The solid lines represent the baseline model and are identical to the results presented in the right panel of Figure 11. The dashed lines represent $C_w^{abilities}(\kappa_{AI})$ when $A_o^{abilities}(\kappa_{AI})$ is used instead of $A_o^{skills}(\kappa_{AI})$.

O*NET defines the work activities as "general types of job behaviors occurring in multiple jobs."[33] The O*NET website categorises these activities into four groups: i) information input, ii) interacting with others, iii) mental processes, and iv) work output. These categories encompass a variety of tasks, from "Coordinating the Work and Activities of Others" and "Developing

[32]Note that we keep our baseline approach to identify core and side variables: the top 33% of cognitive abilities/work activities are identified as "core" in each occupation.

[33]see `https://www.onetonline.org/find/descriptor/browse/4.A` The website was accessed in June 2024.

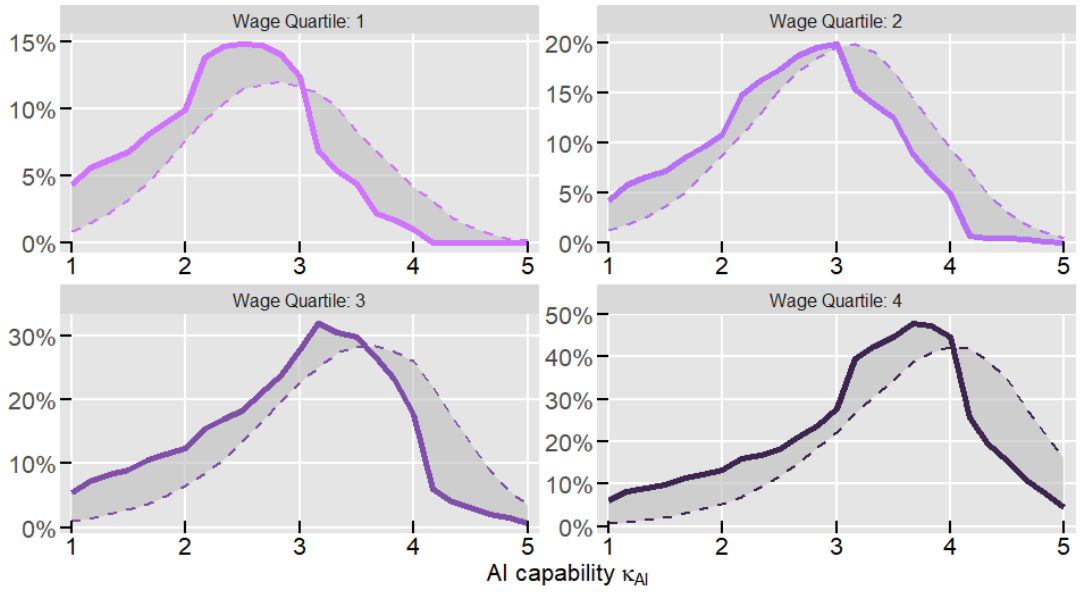
Objectives and Strategies" to "Identifying Objects, Actions, and Events" and "Interpreting the Meaning of Information for Others". Although there is significant variation among these work activities, we exclude "work output" only due to its focus on physical activities, which are not relevant to AI interactions in our scope. We end up with 41 different work activities. The *difficulty* level $L_{o,s}^{work\,act.}$ and importance $I_{o,s}^{work\,act.}$ of these activities are defined similarly to the O*NET skills table, but the difficulty level is set with a range of 0 to 7. We rescaled this to a 0 to 6 range to retain compatibility of the results with our baseline model. The importance variable $I_{o,s}^{work\,act.}$ maintains the same range as in the skills table, i.e. from 1 to 5.

We calculate $A_o^{work\,act.}(\kappa_{AI})$ in the same fashion as shown previously, i.e. by utilising the level $L_{o,s}^{work\,act.}$ and importance $I_{o,s}^{work\,act.}$ variables from the work activities tables from O*NET:

$$A_o^{work\,act.}(\kappa_{AI}) = \frac{\sum_{s\in S} A_{o,s}^{work\,act.}(\kappa_{AI}) \cdot I_{o,s}^{work\,act.}}{\sum_{s\in S} I_{o,s}^{work\,act.}}; \quad A_{o,s}^{work\,act.}(\kappa_{AI}) = \begin{cases} 0 & \text{if } L_{o,s}^{work\,act.} > \kappa_{AI}, \\ 1 & \text{if } L_{o,s}^{work\,act.} \leq \kappa_{AI} \end{cases}$$

Figure 16 compares $C_o(\kappa_{AI})$ of the baseline results with the alternative case that uses restricted work activities dataset. The work activities show more variation compared to the results from cognitive abilities, with peak levels averaging 3.3pp lower and occurring around 0.25 κ_{AI} later. Despite these deviations, our findings remain robust to changes in the tables used from the O*NET.

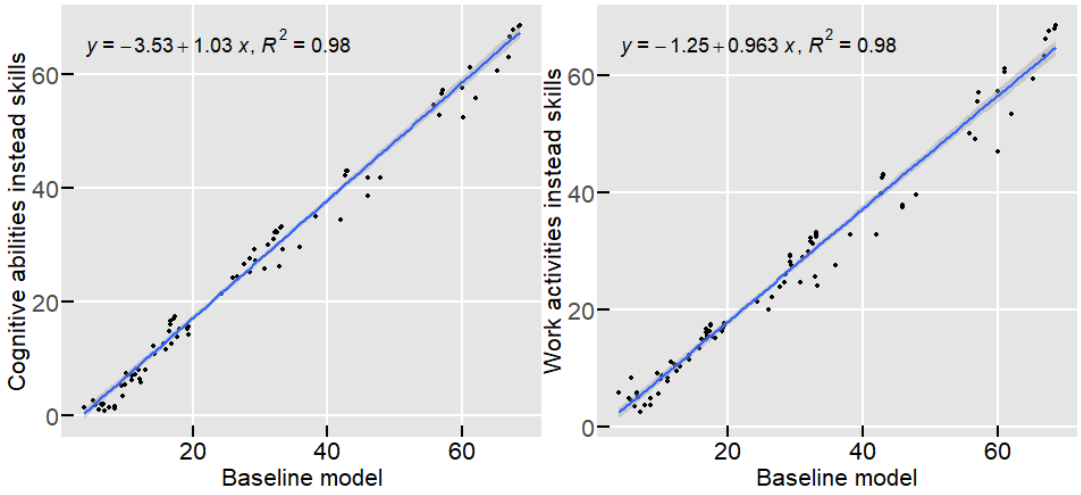
Figure 16: Work activities instead of skills



Note: The figure compares complementarity effects $C_w(\kappa_{AI})$ of the baseline model with the alternative approach that uses work activities table instead of skills. Each panel shows the evolution of $C_w(\kappa_{AI})$ for individual wage quartiles against the progressing AI capabilities κ_{AI}. The solid lines represent the baseline model and are identical to the results presented in Figure 11 right panel. The dashed lines represent $C_w^{work\,act.}(\kappa_{AI})$ when $A_o^{work\,act.}(\kappa_{AI})$ is used instead of $A_o^{skills}(\kappa_{AI})$.

The comparisons of AISA between the baseline model and the alternatives that use either cognitive abilities or work activities instead of skills reveal no substantial differences between the approaches, as shown in Figure 17. Cognitive abilities are only 3.53 percentage points lower on average, and the coefficient of determination R^2 reaches 0.98. The comparison with the work activities provides very similar results.

Figure 17: AISA in comparison to alternative models (skills-oriented)



Note: The figure presents scatter plots of AISA (as used in Figure 7) from the baseline model (x-axis) and AISA calculated with the alternative setups of the baseline model used for the robustness analysis (y-axis). We consider two different setups: i) cognitive abilities table instead of skills table, and ii) work activities table instead of skills table. Panels also display summaries of simple OLS regression lines.

5.3 Using different definitions of core and side skills

As a last robustness exercise, we investigate the robustness of the results with regard to changing the definition for core and side skills. The baseline approach defines core skills as the top 33.3% of the most important skills for an occupation.[34] We construct two alternative models to inspect the impact of different ratios used to identify the core skills on complementarity $C_o(\kappa_{AI})$.

The first model defines core skills as the top 50% based on importance, thus including a broader set of skills. The second model narrows the core skills to the top 20%. Figure 18 compares the baseline results with these alternative approaches on the evolution of complementarity $C_o(\kappa_{AI})$. When 50% of skills are considered core (dotted lines), $C_o(\kappa_{AI})$ peaks at a slightly lower AI capability (on average by 0.5) and is more pronounced, with a maximum difference of up to 5 percentage points. Conversely, when only 20% of the most important skills are considered core (dashed lines), the dynamics reverse: $C_o(\kappa_{AI})$ improves more gradually, and the maximum complementarity is slightly reduced compared to the baseline model.

[34]This approach mimics O*NET methodology to define core tasks (see subsection 3.2)

Figure 18: Different core skills identification process



Note: The figure compares complementarity effects $C_w(\kappa_{AI})$ of the baseline model with the two alternative specifications of core skills. The baseline model considers core skills to be the top 33.3% of the most important skills for an occupation. The first alternative model defines core skills as the top 50% most important skills (broader definition of core skills) and is presented with dotted lines. The second alternative model defines core skills as the top 20% most important skills (a more narrow definition of core skills) and is presented with dashed lines. The solid lines represent the baseline model and are identical to the results presented in Figure 11 right panel. Each panel shows the evolution of $C_w(\kappa_{AI})$ for an individual wage quartile against the progressing AI capabilities κ_{AI}.

6 Conclusion

In this paper, we have explored the impact of AI's evolving capabilities on the labour market, focusing on the exposure of 711 US occupational categories and the potential of AI to complement or substitute human labour. We model how the share of an occupation's skills exposed to AI depends on the difficulty of these skills and the AI's cognitive capability. Our analysis distinguishes between the impacts on core and side skills and investigates differential exposures across wage quartiles.

We first investigate the overall exposure of industries across occupations and the wage distribution. We find that AI may initially affect occupations uniformly across the wage spectrum, impacting approximately 7% of skills at lower AI capability levels. However, as capabilities improve, up to 45% of skills in the highest wage quartile are susceptible to automation by AI, compared to only 26% in the lowest quartile.

Nevertheless, looking into the impact on core and side skills, we find that AI may still lead to increasing inequality as it will tend to substitute low-wage work more easily than high-wage work. We find that low AI capabilities complement all workers, as side skills are simpler than core skills. However, as AI capabilities advance, core skills in lower-wage jobs become exposed, threatening substitution and increased inequality. In contrast to the intuitive notion that the

rise of AI may harm white-collar workers, we find that those remain safe as their core skills are hard to automate. Yet workers will be forced to adapt to these heavily AI-supported new working modalities. They will need to learn to trust AI-enabled services supporting them in their side tasks while focusing more of their work's attention on their core tasks.

The simulations in this paper provide a detailed view of who will be impacted by AI and in what ways, offering valuable insights for economic policy formulation. As AI continues to reshape the labour market, several considerations need careful evaluation. Enhancing skill development and training is crucial for workers at risk of AI-driven displacement, which can be focused on adapting workforce capabilities in the affected skill areas and occupations. Transparency in AI deployment and involving workers in implementation decisions are essential to ensure that AI complements rather than replaces human labour. Additionally, establishing safety nets and transition programmes can support those adversely affected by AI. On a broader scale, strengthening international cooperation on AI labour policies allows for a unified approach to managing AI's impact globally. These considerations are vital for policymakers to balance the benefits and challenges of AI in the workforce.

References

Acemoglu, D. (2024): "The simple macroeconomics of ai." *Economic Policy* .

Acemoglu, D., D. Autor, J. Hazell, & P. Restrepo (2022): "Artificial intelligence and jobs: Evidence from online vacancies." *Journal of Labor Economics* **40(S1)**: pp. S293–S340.

Acemoglu, D. & P. Restrepo (2018a): "Artificial intelligence, automation, and work." In "The economics of artificial intelligence: An agenda," pp. 197–236. University of Chicago Press.

Acemoglu, D. & P. Restrepo (2018b): "Low-skill and high-skill automation." *Journal of Human Capital* **12(2)**: pp. 204–232.

Acemoglu, D. & P. Restrepo (2018c): "The race between man and machine: Implications of technology for growth, factor shares, and employment." *American economic review* **108(6)**: pp. 1488–1542.

Acemoglu, D. & P. Restrepo (2019): "Automation and new tasks: How technology displaces and reinstates labor." *Journal of economic perspectives* **33(2)**: pp. 3–30.

Acemoglu, D. & P. Restrepo (2020): "Robots and jobs: Evidence from us labor markets." *Journal of political economy* **128(6)**: pp. 2188–2244.

Acemoglu, D. & P. Restrepo (2022): "Tasks, automation, and the rise in us wage inequality." *Econometrica* **90(5)**: pp. 1973–2016.

Adachi, D., D. Kawaguchi, & Y. U. Saito (2024): "Robots and employment: Evidence from japan, 1978–2017." *Journal of Labor Economics* **42(2)**: pp. 591–634.

Albanesi, S., A. D. Da Silva, J. F. Jimeno, A. Lamo, & A. Wabitsch (2023): "New technologies and jobs in europe." *Technical report*, National Bureau of Economic Research.

Aldasoro, I., O. Armantier, S. Doerr, L. Gambacorta, & T. Oliviero (2024): "The gen ai gender gap." *Economics Letters* **241**: p. 111814.

Arntz, M., T. Gregory, & U. Zierahn (2016): "The risk of automation for jobs in oecd countries: A comparative analysis." *OECD Social, Employment, and Migration Working Papers* **(189)**: p. 0_1.

Autor, D. (2019): "Work of the past, work of the future." *Technical report*, National Bureau of Economic Research.

Autor, D. H. (2015): "Why are there still so many jobs? the history and future of workplace automation." *Journal of Economic Perspectives* **29(3)**: pp. 3–30.

Autor, D. H. & D. Dorn (2013): "The growth of low-skill service jobs and the polarization of the us labor market." *American economic review* **103(5)**: pp. 1553–1597.

Bessen, J. E. (2016): "How computer automation affects occupations: Technology, jobs, and skills." *Boston Univ. school of law, law and economics research paper* **(15-49)**.

Brynjolfsson, E., D. Li, & L. R. Raymond (2023): "Generative ai at work." *Technical report*, National Bureau of Economic Research.

Brynjolfsson, E., T. Mitchell, & D. Rock (2018): "What can machines learn and what does it mean for occupations and the economy?" In "AEA papers and proceedings," volume 108, pp. 43–47. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Cazzaniga, M., F. Jaumotte, L. Li, G. Melina, A. J. Panton, C. Pizzinelli, E. J. Rockall, & M. M. Tavares (2024): "Gen-ai: Artificial intelligence and the future of work." *Staff Discussion Notes* **2024(001)**.

Cheng, Z., D. Lee, & P. Tambe (2022): "Innovae: Generative ai for understanding patents and innovation." *Available at SSRN 3868599* .

Cockburn, I. M., R. Henderson, & S. Stern (2018): *The impact of artificial intelligence on innovation*, volume 24449. National bureau of economic research Cambridge, MA, USA.

Dekle, R. (2020): "Robots and industrial labor: Evidence from japan." *Journal of the Japanese and International Economies* **58**: p. 101108.

Eisfeldt, A. L., G. Schubert, & M. B. Zhang (2023): "Generative ai and firm values." *Technical report*, National Bureau of Economic Research.

Eloundou, T., S. Manning, P. Mishkin, & D. Rock (2023): "Gpts are gpts: An early look at the labor market impact potential of large language models." *arXiv preprint arXiv:2303.10130* .

Felten, E., M. Raj, & R. Seamans (2021): "Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses." *Strategic Management Journal* **42(12)**: pp. 2195–2217.

Frey, C. B. & M. A. Osborne (2017): "The future of employment: How susceptible are jobs to computerisation?" *Technological forecasting and social change* **114**: pp. 254–280.

Furman, J. & R. Seamans (2019): "Ai and the economy." *Innovation policy and the economy* **19(1)**: pp. 161–191.

Gmyrek, P., J. Berg, & D. Bescond (2023): "Generative ai and jobs: A global analysis of potential effects on job quantity and quality." *ILO Working Paper* **96**.

Hatzius, J. *et al.* (2023): "The potentially large effects of artificial intelligence on economic growth (briggs/kodnani)." *Goldman Sachs* .

Korinek, A. & M. Juelfs (2022): "Preparing for the (non-existent?) future of work." *Oxford Handbook of AI Governance* .

Michaels, G., A. Natraj, & J. Van Reenen (2014): "Has ict polarized skill demand? evidence from eleven countries over twenty-five years." *Review of Economics and Statistics* **96(1)**: pp. 60–77.

Noy, S. & W. Zhang (2023): "Experimental evidence on the productivity effects of generative artificial intelligence." *Science* **381(6654)**: pp. 187–192.

Peng, S., E. Kalliamvakou, P. Cihon, & M. Demirer (2023): "The impact of ai on developer productivity: Evidence from github copilot." *arXiv preprint arXiv:2302.06590* .

Perez-Cruz, F. & H. S. Shin (2024): "Testing the cognitive limits of large language models." *Technical report*, Bank for International Settlements.

Pizzinelli, C., A. J. Panton, M. M. M. Tavares, M. Cazzaniga, & L. Li (2023): *Labor market exposure to AI: Cross-country differences and distributional implications*. International Monetary Fund.

Svanberg, M., W. Li, M. Fleming, B. Goehring, & N. Thompson (2024): "Beyond AI exposure: Which tasks are cost-effective to automate with computer vision?" *Working paper*, MIT.

Tolan, S., A. Pesole, F. Martínez-Plumed, E. Fernández-Macías, J. Hernández-Orallo, & E. Gómez (2021): "Measuring the occupational impact of ai: tasks, cognitive abilities and ai benchmarks." *Journal of Artificial Intelligence Research* **71**: pp. 191–236.

Webb, M. (2020): "The impact of artificial intelligence on the labor market." *Available at SSRN 3482150* .

Yang, C.-H. (2022): "How artificial intelligence technology affects productivity and employment: firm-level evidence from taiwan." *Research Policy* **51(6)**: p. 104536.

A Appendix

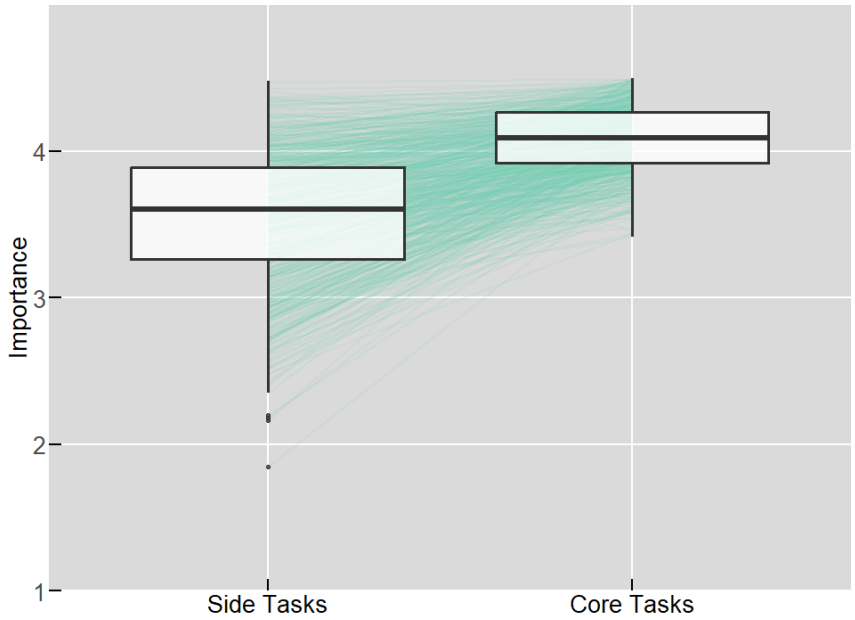
Table A1: Tasks descriptions for the occupation *Fundraising Managers* from O*NET database

No.	Task Description
1	Assign, supervise, and review the activities of fundraising staff.
2	Compile or develop materials to submit to granting or other funding organisations.
3	Conduct research to identify the goals, net worth, charitable donation history, or other data related to potential donors, potential investors, or general donor markets.
4	Contact corporate representatives, government officials, or community leaders to increase awareness of organisational causes, activities, or needs.
5	Design and edit promotional publications, such as brochures.
6	Develop fundraising activity plans that maximise participation or contributions and minimise costs.
7	Develop strategies to encourage new or increased contributions.
8	Direct activities of external agencies, establishments, or departments that develop and implement fundraising strategies and programs.
9	Establish and maintain effective working relationships with clients, government officials, and media representatives and use these relationships to develop new fundraising opportunities.
10	Establish goals for soliciting funds, develop policies for collection and safeguarding of contributions, and coordinate disbursement of funds.
11	Evaluate advertising and promotion programs for compatibility with fundraising efforts.
12	Formulate policies and procedures related to fundraising programs.
13	Manage fundraising budgets.
14	Plan and direct special events for fundraising, such as silent auctions, dances, golf events, or walks.
15	Produce films and other video products, regulate their distribution, and operate film library.
16	Write interesting and effective press releases, prepare information for media kits, and develop and maintain company internet or intranet Web pages.

Table A2: Examples for Skill Levels: Corresponding Occupations from O*NET database

AI capability /Skills	Troubleshooting	Critical Thinking	Active Listening	Mathematics
1.5	Social and Community Service Managers	Cleaners of Vehicles and Equipment	Pressers, Textile, Garment, and Related Materials	Writers and Authors
2.0	General and Operations Managers	-	-	Human Resources Specialists
2.5	Computer and Information Systems Managers	Court Reporters and Simultaneous Captioners	Terrazzo Workers and Finishers	Managers, All Other
3.0	Industrial Production Managers	Farm Labor Contractors	Cooks, Fast Food	Advertising and Promotions Managers
3.5	Industrial Engineers	Property, Real Estate, and Community Association Managers	Architectural and Engineering Managers	Financial Managers
4.0	Electro-Mechanical and Mechatronics Technologists	Advertising and Promotions Managers	Chief Executives	Cost Estimators
4.5	Aircraft Mechanics and Service Technicians	Anesthesiologists	Labor Relations Specialists	Operations Research Analysts
5.0	-	Judges, Magistrate Judges, and Magistrates	Judges, Magistrate Judges, and Magistrates	Mathematicians

Figure A1: Importance of Core and Side Tasks according to O*NET database



Note: The figure supports our proposition that the skills can be split into core and side skills based on their importance as core tasks are more important than the side tasks on occupational average.

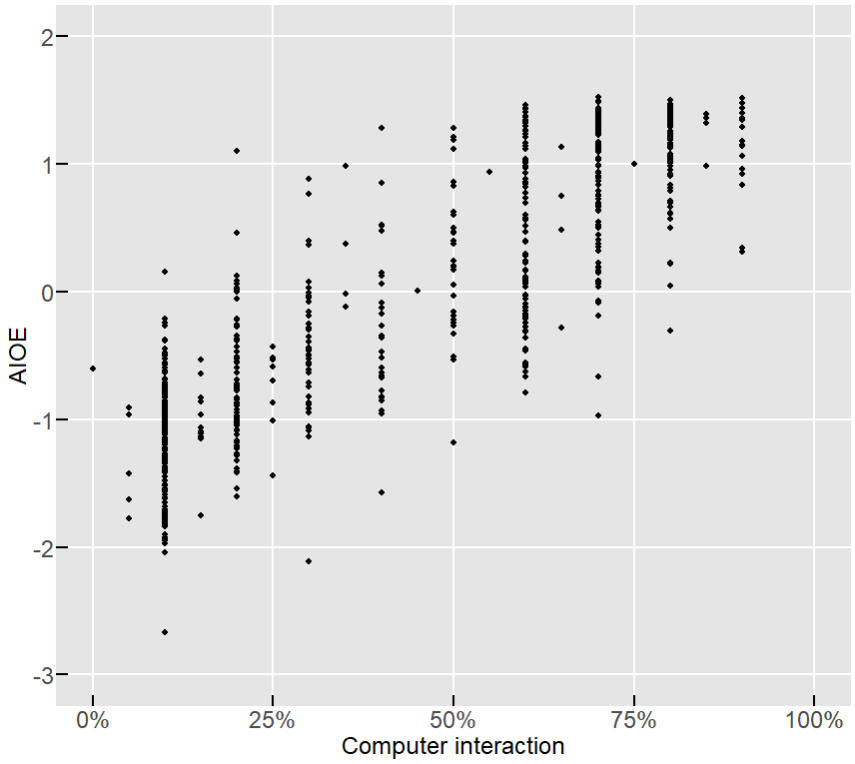
Box A1: Prompt to define fractions of time spent in occupations

"Below is a list of task descriptions for the profession of {occupation}: {tasks}: With this description compile a JSON file with an estimate of how much of the worktime is spent on:

1. *Working on computer*
2. *Talking to people*
3. *Physical activities*

Note: The prompt was looped across all occupations in the dataset. {occupation} and {tasks} represent two variables varying in the prompt.

Figure A2: Comparison between computer interaction and AIOE on individual O*NET occupations



Note: Figure presents the relationship between computer interaction variable and AIOE developed by Felten *et al.* (2021). The correlation between these variables is notably high at 86.3%.