



AALBORG UNIVERSITY
DENMARK

Aalborg Universitet

The impact of artificial intelligence on skills at work in Denmark

Holm, Jacob Rubæk; Lorenz, Edward

Published in:
New Technology, Work and Employment

DOI (link to publication from Publisher):
[10.1111/ntwe.12215](https://doi.org/10.1111/ntwe.12215)

Publication date:
2022

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Holm, J. R., & Lorenz, E. (2022). The impact of artificial intelligence on skills at work in Denmark. *New Technology, Work and Employment*, 37(1), 79-101. Advance online publication. <https://doi.org/10.1111/ntwe.12215>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

This is the peer reviewed version of the following article: Holm, J.R. & Lorenz, E. (2022) The impact of artificial intelligence on skills at work in Denmark. *New Technology, Work and Employment*, 37, 79– 101, which has been published in final form at <https://doi.org/10.1111/ntwe.12215>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions. This article may not be enhanced, enriched or otherwise transformed into a derivative work, without express permission from Wiley or by statutory rights under applicable legislation. Copyright notices must not be removed, obscured or modified. The article must be linked to Wiley's version of record on Wiley Online Library and any embedding, framing or otherwise making available the article or pages thereof by third parties from platforms, services and websites other than Wiley Online Library must be prohibited.

The Impact of Artificial Intelligence on Skills at Work in Denmark

Abstract

Based on a unique dataset on the use of artificial intelligence (AI) among employees in Denmark, we investigate within-job relationships between AI use and skill requirements. We show that the effects of AI are varied and depend on providing orders to humans versus providing information for further human handling and in which occupation it is used. AI may enhance or augment skills through, for example, the increased use of high-performance work practices, or it may simply increase work pace constraints and reduce employee autonomy. The results imply that the diffusion of AI can increase inequalities in the labour market by augmenting skills used in high-skill jobs while having relatively more adverse impacts on other jobs. We use additive noise modelling to establish the likely direction of causality in our results and find that the direction of causality is from AI use to skill requirements.

Keywords: Artificial intelligence, skills, job requirements approach, constraints, autonomy, learning, high-performance work practices, additive noise modelling

Running title (max 40 characters): Artificial Intelligence, Skills and Work

Authors:

Jacob Rubæk Holm (corresponding)

jrh@business.aau.dk

IKE/DRUID, Aalborg University Business School,

Fibigerstræde 11, 9220 Aalborg Ø, Denmark

ORCID: 0003-1110-1918

Short bio for JR Holm:

Jacob Rubæk Holm is Associate Professor of Industrial Dynamics and Quantitative Methods at Aalborg University Business School, Denmark. His research focuses on regional, industrial and structural change over time, ranging from studies of specific industries or specific regions to theoretical and methodological contributions with an evolutionary angle.

Edward Lorenz

lorenz@business.aau.dk

IKE/DRUID, Aalborg University Business School,

Fibigerstræde 11, 9220 Aalborg Ø, Denmark

Short bio for E Lorenz:

Edward Lorenz is a part-time professor at Aalborg University, Denmark and Emeritus Professor of Economics at the University of Cote d'Azur, France. He is a visiting professor at the University of Johannesburg, South Africa. His research focuses on the internationally comparative analysis of business organization, employment relations and innovation systems.

Acknowledgements:

We thank participant at the joint online workshop on “Technological Change and Regional Inequalities” by the School of Business, Economics and Law at the University of Gothenburg and Aalborg University Business School for valuable comments. Any errors are our own.

The Impact of Artificial Intelligence on Skills at Work in Denmark

Abstract

Based on a unique dataset on the use of artificial intelligence (AI) among employees in Denmark, we investigate within-job relationships between AI use and skill requirements. We show that the effects of AI are varied and depend on providing orders to humans versus providing information for further human handling and in which occupation it is used. AI may enhance or augment skills through, for example, the increased use of high-performance work practices, or it may simply increase work pace constraints and reduce employee autonomy. The results imply that the diffusion of AI can increase inequalities in the labour market by augmenting skills used in high-skill jobs while having relatively more adverse impacts on other jobs. We use additive noise modelling to establish the likely direction of causality in our results and find that the direction of causality is from AI use to skill requirements.

Key words

Artificial intelligence, skills, job requirements approach, constraints, autonomy, learning, high-performance work practices, additive noise modelling

1. Introduction

There is an on-going debate on the effects of new and emerging technologies on jobs and skills (Frey and Osborne, 2017; Brynjolfsson et al., 2018; Felten et al., 2019). Much of the literature has focused on robots and artificial intelligence (AI) and has attempted to assess how the adoption of these technologies will affect the employment and skills of different occupational groups. While the literature on robotics has drawn on the results of data collected mainly at the sector and national levels in the 1990s,¹ the research on the effects of AI has been more speculative in nature and has sought to predict future patterns of adoption based on technical assessments of the susceptibility of different detailed occupations to substitution by AI. This paper contributes to this literature with an analysis of the impact of AI on the skills employees use in their daily work based on what we believe to be a unique large-scale, employee-level survey with measures of the use of AI. The survey results allow us to identify two basic types of AI used in daily work, distinguishing between tasks where the employee receives orders or directions generated automatically by a computer or computerized machinery and tasks where the employee makes use of information compiled automatically by a computer or computerized machinery for further decision making or for advising clients or customers. We thus focus on AI defined as computer systems performing tasks that normally require human intelligence and on how such systems are used to automate and augment tasks in jobs. This focus may be contrasted both with ‘true’ AI that is capable of replicating and outperforming human intelligence in all respects (Agraval et al., 2019: 3), and with AI defined more narrowly as a machine learning (ML) prediction technology. Our way of measuring AI encompasses but is not limited to ML. We find that AI

¹ There is a large body of literature that draws mainly on the data collected by the International Federation of Robotics (IFR) on the installation of industrial robots dating from the 1990s. See in particular Acemoglu and Restrepo (2017); Dauth et al. (2017); Graetz and Michaels (2018) and Giuntella and Wang (2019).

most often augments skills used in high-skill jobs, while other jobs more often become more constrained in terms of work pace.

The paper is structured in the following manner. In Section 2, we present an overview of the recent economics literature on the adoption of AI and its impact on employment and skills. We point to some of the limitations of this literature in terms of capturing how AI may transform employees' skills in daily work activity as opposed to simply substituting for their skills. In Section 3, we describe the survey design and present descriptive statistics on the adoption of the two forms of AI use we identify according to sector and broad occupational category. In Section 4, our econometric analysis shows that the two different forms of AI use have different impacts on employees' skills at work and that these impacts may differ across high, middling and low skilled occupational categories. Using additive noise modelling (ANM), we provide indications that the direction of causality is from the use of AI at work to skill requirements. In Section 5, we conclude by pointing to the importance of our results for education and training systems and argue that there is a clear need for further micro-level studies on AI and skills.

2. Background to the debate on the impact of AI on employment and skills

Concerns about the impact of automation technologies on employment and skills are far from new. Keynes (1930) is famously credited with popularising the term 'technological unemployment', referring to technology destroying jobs faster than we can discover new ones. In an article focusing on the recent debate on the impact of new technology on employment, Autor (2015) notes that during the 1950s and 1960s concern in the US led the Johnson administration to set up a commission on automation and employment. Braverman's (1974) book on the labour process gave rise to a large debate during the 1970s and 1980s on whether management has an interest in pursuing a general strategy of deskilling to increase managerial control. Braverman distinguished between organisational deskilling and technological deskilling and with respect to the latter identified the use of numerical control (NC) machine tools as an example of how management may use automation technologies to increase their control by separating the tasks of conception from execution.² The 1990s witnessed a burgeoning literature in economics on skill-biased technological change (Katz and Murphy, 1992; Machin and van Reenen (1998) which attributed the secular decline in the relative wages of low skilled workers to the effects of technological change.³ This literature argued that technology substitutes for the skills of workers in low paid jobs and augments the skills of the higher paid occupations. A related strand of literature focused on the skill-biased effects of organizational change, arguing that organizational changes such as the delayering of hierarchies, which enhances worker autonomy, can have independent effects on the occupational distribution by reducing demand for workers with the lowest skills. (Caroli and Van Reenen 2001; Piva et al., 2005).

2.1. The task-based approach

² For a useful discussion of Braverman's thesis of deskilling and the subsequent research that qualified or criticized his views, see Noon et al. (2013, pp. 147-56).

³ For a survey of the literature on skill-biased technological change, see Sanders, M. and ter Weel, B., (2000).

The recent literature on the effects AI on skills has positioned itself in relation to the research on the skill bias of technical change. A striking feature of this literature is that it has investigated the impact of AI on employment and skills without any attempt to measure what takes place inside the enterprise. The 24 contributions to a recent NBER volume edited by Agravall et al. (2019) focus on the implications of AI for employment, wages and economic growth, but none deal explicitly with effects at the level of firms, employees or jobs. The possible exception is the contribution by Athey (2019), which can be seen as an argument that AI will change the jobs of research economists by automating some aspects of econometrics. The lack of focus on the workplace is a defining feature of the task-based approach to studying the impact of new technologies on jobs and skills, associated notably with the seminal work of Autor et al. (2003) on the effects of computerization on labour market polarization beginning in the 1980s. The task-based approach uses the standardized descriptors of the different mix of knowledge, skills and abilities required for detailed occupations provided in the O*NET or its precursor the US Dictionary of Occupational Titles (DOT) as a basis for interpreting medium-term patterns of change in occupational shares in the labour market. The O*NET descriptors are compiled in a way that draws on employee-level information, as they are based on a measurement program that involves interviewing employees using standardized questionnaires from a random sample of US businesses.⁴ However, the objective of the O*NET is not to identify any differences that may exist amongst employees in the same occupational category but rather to provide useful and general information to job seekers and employers about the types of skills and training required for different career options.

The work by Autor et al. (2003) and researchers such as Goos and Manning (2007) on labour market polarization has generated important new insights into the effect of computers on employment and skills. As discussed in Holm et al. (2020), the task-based approach they use is nonetheless limited by its reliance on aggregate data at the sector and occupational levels that precludes investigating within-firm effects related to differences in investments in new technology and in the adoption of organizational practices. Consequently, the contribution of within-firm effects to the observed patterns of change in employment and skills remains unexplained. The limitation of investigating the effect of new automation technologies on employment and skills without the insights derived from micro-level data has arguably become more evident in the debate engendered by research by Frey and Osborne (2017) on the future effects of AI in the form of ML on employment. As is well known, Frey and Osborn (2017) came up with the alarmist prediction that 47% of people currently working in the US are at high risk (70% chance or greater) of having their jobs automated over the coming decade or so. Their methodology follows that of the earlier work by Autor and others on job polarization in using the O*NET as a basis for characterizing the skills and knowledge requirements for the tasks of detailed occupational categories across the economy. ML researchers from the University of Oxford assessed whether occupations were susceptible to automation with AI using 70 handpicked occupational descriptors from the O*NET. These assessments considered various technical bottlenecks to fully automating the occupational tasks given the current state of AI technology. The assessments served as the training data for a Gaussian process classifier, which was then used to estimate the probability of automation for all 702 occupations included in the O*NET database (Frey and Osborne, 2017).

⁴ For a description of the measurement program, see: <https://www.onetcenter.org/dataCollection.html>.

2.2 AI impacts and the limitations of the task-based approach

Although creative, the methodology is arguably flawed in several respects. As also discussed by Lloyd and Payne (2019), one limitation has to do with the assumption that all employees with the same occupational category within a sector and across countries are identical in terms of their tasks and skills. The prediction bias this assumption may engender is discussed in a study by Arntz et al. (2016) focusing on the set of 21 OECD countries covered in the first wave of the PIAAC survey, an employee-level survey of adult skills.⁵ Although Arntz et al. (2016) use the same at-risk-of-automation assessments for the 70 handpicked occupational task descriptors used in Frey and Osborne (2017), their study differs by using an imputation procedure to map these at-risk assessments onto employee-level data collected via the background questionnaire in the PIAAC survey measuring skills requirements for individual workers.⁶ Their analysis arrives at a substantially lower estimate of the share of workers at high risk of automation, ranging from 6% for Korea and Estonia to a high of about 13% for Austria and Germany. In the case of the US, they estimate that only about 9% are at high risk.

A second limitation of the approach used by Frey and Osborne (2017) to predict the impact of AI that is shared by Arntz et al. (2016) relates to their use of estimates of the scientific susceptibility of tasks to automation. This fails to consider that the adoption of AI may be constrained by the possibly costly reorganization of work necessary to unbundle those tasks in a job or occupation that can be automated from those that cannot. A recent study by Brynjolfsson et al. (2018), while still using the O*NET to describe the skills and knowledge needed for detailed occupations, addresses the issue of unbundling. Brynjolfsson et al. (2018) develop a 24-item rubric designed to capture the extent to which work activity is well structured in the sense that there is a clear mapping between inputs or actions and outputs that can be learnt by a ML neural net with sufficient data. If there are difficulties in measuring this relationship, as is the case for unstructured social interaction or a lot of complex cognitive work, then the work activity receives a low susceptible to ML (SML) score. Using CrowdFlower, a human intelligence task crowd sourcing platform, they apply the rubric to 2,056 direct work activities shared across occupations they identify based on the occupational task descriptors in the O*NET.⁷

Brynjolfsson et al. (2018) find in general that there is considerable variability across occupations in the susceptibility of their component tasks to automation, with only a few having high SML scores for all tasks. This implies that automation technologies are unlikely to result in the elimination of entire occupations. The authors conclude that significant job redesign will be needed ‘for unleashing ML potential’ and that ‘The focus of researchers, as well as managers and entrepreneurs, should be not (just) on automation, but on job redesign’ (Brynjolfsson et al., 2018: 44).

⁵ PIAAC: The Programme for the International Assessment of Adult Competencies.

⁶ The OECD’s Survey of Adult Skills (PIAAC) uses a household survey frame. The main objective of the survey is to measure adult literacy, numeracy and ICT-related problem-solving skills. The survey includes a background questionnaire with a module addressed to employees that is designed to measure adult skills at work using the job requirements approach (JRA) developed for the UK Skills Surveys (Felstead et al., 2002). Rather than measuring skills by asking employees for a subjective assessment of what skills they have, the JRA measures skills by asking employed persons what they do at work. In this way, the survey design seeks to avoid biases associated with the tendency of employees to be overconfident about the level of their own skills. For a discussion of the conceptual framework of the JRA, see OECD (2009).

⁷ For details on the methodology, see Brynjolfsson et al. (2018). For the detailed rubric, see www.sciencemag.org/content/358/6370/1530/suppl/DC1.

A third limitation common to all three studies cited above is the exclusive focus on the future substitution effects of AI and the failure to investigate how AI may complement or augment the skills of existing occupations as opposed to replacing them. The exclusive focus of the paper by Frey and Osborne (2017) on the predicted substitution effects may in part be justified by the fact that their research was carried out in 2013 at the beginning of a new wave of adoption of AI in industry resulting from advances in ML methods. Our knowledge of ML methods was understandably limited at this point in time, and Frey and Osborne were not alone in predicting that ML was the harbinger of a technological singularity that would progressively eliminate all or most human labour in industrial production.⁸ Today, more than seven years after the original Frey and Osborne (2013) working paper appeared, there has been a considerable diffusion of AI, including ML methods, in industry that potentially provides the empirical basis for assessing the extent to which it substitutes for or complements human labour. However, attempts to investigate the possibly complementary effects of AI on skills have been hampered by the task-based approach's reliance on aggregate employment data and the O*NET as a basis for investigating AI impacts. This precludes an investigation of possibly heterogeneity in the within-job changes in skills that may result from an individual employee's use of AI. The importance of these within-job effects can only be assessed with micro-level data at the employee level.

This limitation of the task-based approach in this respect is illustrated by a recent paper by Felten et al. (2019). To estimate task susceptibility to AI automation, the authors draw on the Electronic Frontier Foundation (EFF) AI Progress Measurement dataset that tracks advances since 2010 in specific applications of AI, such as image recognition, translation or the ability to play strategic games. These measurements are mapped onto the occupational descriptors in the O*NET to construct AI occupational impact (AIOI) measures for occupations at the six-digit Standard Occupational Classification (SOC) level. As the authors observe, these impact measures cannot distinguish between impacts that are substituting for as opposed to complementing human labour, and in this respect the analysis makes little headway in analysing what the impact of AI on skills has been. To infer what this impact may have been, the authors perform a regression analysis at the occupational level to estimate the relationship between the AI impact score and the observed change in occupational employment and wages for the US using the data from the Bureau of Labor Statistics (BLS) for each occupation from 2010 to 2016. They conclude that while AI has not had an impact on employment in the US, it has had a small but positive impact on wages, and it is the high-income groups that have benefited the most. They conclude by arguing that their findings suggest that complementary skills effects are important and that by increasing the relative earnings of the high-skills groups AI may have increased labour market polarization.

Although the paper has the merit of seeking to analyse the impact of AI based on the actual changes in wages and employment for occupations that have occurred between 2010 and 2016, it is limited for

⁸ See in particular Brynjolfsson and McAfee (2014). For a critical discussion of the singularity hypothesis, see Nordhaus (2015). Also see Aghion et al. (2019), who analyze different potentials for a singularity in a growth model. They argue that Baumol's cost disease, that is, the presence of economic activities that are '*essential but hard to improve*', prohibits a singularity. Another result from their model is that if AI is used in the innovation process AI will replicate innovation instantly, thus removing incentives from innovation and prohibiting a singularity by acting as a break on innovation.

one of the reasons discussed above. It assumes that all employees with the same detailed occupational title are affected in the same way, and this abstracts from an analysis of within-job changes in skills that is the natural level of analysis to adopt if one wants to investigate the way AI may complement the skills of employees whose jobs are transformed but not eliminated by AI. In this paper, we present the results of what we believe to be the first study of the impacts of AI on skills using micro data at the employee level: the TASK survey carried out at Aalborg University Business School in 2019. While the cross-sectional survey design does not allow us at present to investigate the employment impacts of AI, we are able to investigate the within-job effects of AI on skills at work. Section 3 below presents the survey design and provides descriptive statistics on the diffusion of AI in the Danish economy.

3. The TASK Survey

The survey on technology and skills (TASK survey) in Denmark was carried out in the spring of 2019 by Statistics Denmark. It relied on a stratified sample of employees in workplaces outside public sector administration with at least five full-time equivalent employees and. The questionnaire for TASK was inspired by other major surveys on employees' tasks at work for comparability and includes a number of novel questions on the use of technology at work. Further details and descriptive statistics for the TASK dataset can be found in Gjerding et al. (2020). Some of the early results from the TASK data concerned the differences and similarities between AI and robotics. It was shown that while robotics is used more intensively in some industries than in others, AI is broadly diffused and already in 2019 more than 1 in 4 employees in Denmark use AI with regular intervals compared to less than 1 in 10 for robots (Holm et al., 2021).⁹

In this paper, we are interested in the relationship between the use of AI and the skills required for a job as revealed by the tasks performed in the job. As we are interested in assessing how jobs may be transformed by AI, we set a high bar for using AI and only focus on the effect of using AI on employees who use it on a daily basis. Of employees, 12.56% 'make use of information compiled automatically by a computer or by computerized machinery for making decisions or for advising clients or customers' daily, while 7.90 % 'receive orders or directions generated automatically by a computer or by computerized machinery' daily. We focus on these two types of use and refer to them as 'AI for decisions' and 'AI orders', respectively. These two types of use contrast AI as a tool used for potentially complex and autonomous decision making with AI supplying orders or directions and thus potentially putting constraints on work activity. The TASK data contain an array of variables capturing skill requirements as revealed by the tasks of a job to illuminate these hypothesized relationships between AI use and skills. These variables are described in Section 4. In Section 3.1, we present some descriptive statistics on AI use in jobs.

3.1 Descriptive statistics

⁹ An important proviso regarding the TASK survey data is that while it measures the use of AI and provides the basis for an investigation of the relationship of AI use to job skills, it does not allow for an investigation of the implementation process of AI, and correspondingly the data cannot be used to assess questions relating to whether the observed impacts of AI on skills are those that were intended by management. Nor can it explore possible reconstitution effects linked to the use of AI over time in the firm. For a discussion of the distinction between intended and unintended effects and reconstitution effects, see Edwards and Ramirez (2016).

In the literature on job polarization, occupations are commonly grouped into high-, mid- and low-skill jobs based on the first digit of the ISCO¹⁰ code for the occupation. Table 1 below shows this grouping. If the first digit for an occupation is 1 then the occupation is a managerial occupation, which is classified as high skill. Of employees, 5.59% have a job in this category. The polarization literature generally finds that the proportions of high- and low-skill jobs are increasing, while the proportion of mid-skill jobs is decreasing. This pattern was confirmed for Denmark in Holm et al. (2020).

[Table 1 here]

Table 1 also contains examples of the jobs placed in each category that are observed in our data. As can be seen, the groupings are necessarily somewhat broad, and while they contain jobs of arguably similar skill level, significant qualitative differences are observed within each group. The analysis presented in this paper does not concern the relationship between the use of AI and job polarization—that is, shifts in occupational shares—but rather focuses on changes at the job level associated with AI use. That is, we analyse the effect of AI on the skill requirements of jobs. There are two reasons why AI affects jobs differently; some job types use AI more frequently than others, and using AI can have different effects on different jobs. Figure 1 shows the propensity of the various occupational groups in Table 1 to use AI daily. The reference lines are the national averages presented above and a 45-degree line.

Most occupational groups are above the 45-degree line, meaning that it is more common to use AI daily for decision making than it is to use AI orders daily. The exceptions are groups 8 and 9 (operators and elementary occupations) and the total average for the low-skill group.

[Figure 1 here]

Mid-skill workers are clearly the group that most often use AI, and this is consistent with mid-skill jobs being the group of jobs most exposed to automation technologies in general. However, there is large variation within the group of mid-skill jobs, as group 4 (clerks) comprises the workers most likely to use AI for decisions daily, while group 8 (operators) comprises the workers most likely to use AI orders daily.¹¹ As a whole, the high-skill group has less than average likelihood of using AI orders but average likelihood of using AI for decisions, which sets jobs in this category apart from the other occupational groups.

4. The impact of AI on skills in work

4.1 Measuring job skills

To measure employees’ skills, we adopt the job requirement approach (JRA) used in the OECD’s PIAAC survey referred to above and in other surveys, including the British Skills Survey. The JRA measures the skills employees use in their work by asking about the tasks they perform. For example, if an employee works on a team then this is seen as reflecting the use of team work skills. This means we focus on ‘job skills’ and not on the employee’s ‘own skills’ in the sense of what skills the employee has. As noted in

¹⁰ ISCO: International Standard Classification of Occupations.

¹¹ Despite this difference within the mid group, we use the high-mid-low grouping in the regressions. Using instead the eight groups in Table 1 produced very similar results. The main difference is that as some groups are relatively small, the associated estimates become imprecise when using eight groups.

the methodological report for PIAAC, there may be a discrepancy between the two, and this means that at any point in time a skills mismatch is a possibility. There is also a risk of bias in the sense that the individual respondent might exaggerate what he or she does to present their work in a favourable light. However, the risk of a bias in this sense is likely to be less than it would be if the respondent were asked how good he or she is at performing a task (OECD, 2009: 13).

Table 2 below presents the 14 binary indicators we use to construct measures of job skills in the analysis of AI impacts.¹² These measures are based on those used in successive rounds of the European Working Conditions Survey (EWCS) carried out by Eurofound since the 1990s.¹³ They are designed to identify a set of generic skills that may characterize work activity across different occupations and sectors of the economy. Work on the trend of the skill bias of technical change since the 1980s argues that work has evolved under the impact of information and communication technologies (ICTs) to become more cognitively demanding and to require increased interpersonal skills (Autor and Price, 2013; Katz and Autor, 1999: 1509-38). Our indicators are chosen in part to assess how AI affects these features of work by capturing whether or not an employee's work activity is cognitively demanding, requiring the skills and judgement needed to cope with complex tasks and with learning as opposed to being repetitive and monotonous and whether it requires the skills needed to work autonomously.

We include four indicators that capture whether or not the employee has the social interaction, adaptability and technical skills needed for working in a high-performance work system (HPWS), including the skills needed to work on teams, to work at multiple and flexible tasks and to undertake tasks involving quality control and meeting precise quality standards. There is a variety of research on the characteristics of work organization demonstrating the importance of these tasks and the skills associated with them in both the manufacturing and services sectors. The results of the survey also confirm this for Denmark, as shown by the frequencies of their use in Table 2 below.¹⁴ For example, about 86% of our sample of employees has responsibility for quality control, and about 58% work on teams. We also include indicators of whether the job requires that the employee adapt his or her pace of work to various constraints, including those imposed by the automatic movement of equipment or machinery, the demands of one's boss, production norms or targets set by management and the pace at which one's colleagues work. A striking feature of employees' responses to the survey questions on work pace constraints is how infrequent the former two types of constraints are. Less than 13% of the employees surveyed in Denmark responded that their work pace is constrained by their boss and less than 16% that it is constrained by the automatic movement of equipment or machinery.

It is important to recognize both the merits and shortcomings of the survey-based measures we use for capturing job skills. The questions we use are posed in a simple and for the most part objective manner to increase their reliability in the sense of being unaffected by differences in the type of respondent in terms of occupational category or sector of activity.¹⁵ For example, workers are not

¹² For the survey questions used to measure the 14 indicators, see the Appendix.

¹³ For the objectives of the EWCS and changes in the questionnaires used in successive rounds of the survey, see <https://www.eurofound.europa.eu/surveys/european-working-conditions-surveys-ewcs>.

¹⁴ Classic studies on the core set of work practices used in a HPWS include Appelbaum et al. (2000), Black and Lynch (2004) and Freeman and Kleiner (2000). For evidence on the use of high-performance work practices in the member states of the EU, see Lorenz and Potter (2019) and Holm and Lorenz (2015).

¹⁵ For the extensive cognitive testing of the questions used in the 2010 round of the EWCS designed to assure high reliability and high content validity in the sense of measuring what was intended, see https://www.eurofound.europa.eu/sites/default/files/ef_files/surveys/ewcs/2010/documents/pretest.pdf.

asked whether they exercise autonomy at work, which would likely be interpreted in various ways, but rather are asked if they are able to change or modify their methods or pace of work. This allows us to capture relevant features of work activity and has the advantage of allowing us to assess whether or not these same features are present for employees across the entire sample population. However, this advantage of survey-based methods necessarily comes at the cost of an inability to capture qualitative differences in work activity that may be highly specific to the context in which the individual employee works.

[Table 2 here]

To identify a set of underlying skill domains, we perform a principal component analysis (PCA) on the 14 binary indicators of job skills. We retain the first five components from the PCA each of which has an eigenvalue over 1 and which together account for 54% of the total variance in the dataset.¹⁶ Table 3 shows the factor loadings of the five retained components on the original variables after varimax rotation. Factor 1 loads positively on repetitiveness and monotony and negatively on learning and complexity. We refer to it as Monotony. Factor 2 loads positively on three of the work pace constraint indicators—hierarchical, norm-based and automatic constraints. We refer to it as Constraints. Factor 3 loads positively on two of the work practices used in HPWS, teamwork and job rotation and therefore unsurprisingly also loads positively on horizontal constraints in the sense of work pace being constrained by the work of one’s colleagues. We refer to it as HPWP1. Factor 4 loads positively on exercising control over one’s work methods and work pace. We refer to it as Autonomy. Factor 5 loads positively on two other HPWP—individual responsibility for quality control and meeting precise quality norms or standards. We refer to it as HPWP2.

[Table 3 here]

4.2 Regressions

In this section, we use regression analysis to determine the association between the skill requirements of jobs and AI use. Skill requirements are measured using the five factors or components described above. We regress each of the factors on the two binary indicators for AI use described above. *AI.decisions* takes the value 1 if the employee uses AI for decision making daily and zero otherwise. *AI.orders* takes the value 1 if the employee follows orders generated by AI daily. We use a number of categorical control variables that, except for the variables for occupation, were added to the dataset by Statistics Denmark based on their registry data. Reference categories are in bold. We control for occupation by skill level (**High**, Middling, Low), education (**primary**, secondary or <= 2 years tertiary, >2 years tertiary), Age (**18–39**, 40–59, 60+) and industry (**agriculture**, manufacturing, construction, trade, ICT, finance and real estate, business services, health and education, culture and other services).

By indexing the 14 control variables by j so that z_j is the j th variable the regression, the equation for Factor X becomes

$$FactorX_i = \beta_0 + \alpha_1 AI.decisions_i + \alpha_2 AI.orders_i + \sum_{j=1}^{14} \beta_j z_{j,i} + \epsilon_i, \quad (1)$$

where epsilon is the normally distributed error term. The five regressions following equation 1 are estimated separately with ordinary least squares using the post stratification weights supplied by

¹⁶ For the values of the eigenvalues, see Table 9 in the Appendix.

Statistics Denmark for the TASK survey. Next, we add four interaction terms to the regression equation multiplying the categorical variables for occupations and the indicators of AI use. These interactions will allow us to determine whether the relationship between AI use and skill requirements vary across the three skill groups.

Table 4 shows the result of estimating equation 1 for each of the five factors. Daily use of AI for decision making is found to be positively associated with both HPWP1 and HPWP2 in the job and with constraints on work pace, although the latter is only marginally significant. This means that AI for decision making is associated with increased teamwork, task rotation, quality control, meeting quality standards and constraints on the work pace. Thus, AI for decision making both increases the need for the skills related to adaptability, social interaction and judgement but at the same time tends to fix the pace of work. This result shows that an increase in the use of HPWS practices, such as teamwork or job rotation, does necessarily imply greater discretion in work in the sense of fewer constraints on the work pace.

Relying daily on orders generated by AI, however, only has a relationship with work pace constraints, and this positive relationship is both numerically much larger and statistically more significant than the relationship between AI for decision making and constraints. The estimated magnitudes of the variations in skill requirements related to a synthetic factor are therefore not readily interpretable, but it can be noticed that the magnitudes are at about the same level as the differences by occupation or industry.

The control variables included in the regression model perform largely as expected. After controlling for occupation, there are few effects of industry and of education. However, there are a number of statistically significant effects of age, implying that experience matters for skills at work. The results show that the jobs of older and more experienced workers are on average less monotonous, involve more learning and entail fewer work pace constraints compared to the work of younger colleagues while at the same time being less likely to involve teamwork and task rotation (HPWP1) and to have work pace be constrained by the work of colleagues. However, the work of older and experienced workers more often involves responsibility for quality control and meeting precise quality norms or standards (HPWP2). These differences suggest a process of upskilling that may represent rewards for experience; as workers become more experienced, their work activity tends to involve more learning, they are entrusted with greater responsibility for quality control and their work pace is less constrained. These difference in the results for work experience in relation to HPWP1 and HPWP2 point to the usefulness of using PCA to separate out different underlying skills domains in the data.

[Table 4 here]

In the next set of regressions, we add four interaction terms to each model. The full results are presented in the appendix. Tables 5 and 6 show the resulting differences in the relationship between AI use and skill requirements by occupational group.

[Table 5 here]

AI use was not found to affect monotony and learning in the first set of regressions, but when including the interaction terms it can be seen that daily use of AI for decisions is associated with lower monotony and more learning for high-skill workers, while the daily receipt of orders generated by AI is associated with more monotony and less learning for mid-skill workers.

[Table 6 here]

As seen in Table 5, the positive relationship between AI for decisions and work pace constraints identified earlier is found to only pertain to high-skill workers when adding the interactions, while the positive relationship between AI orders and constraints is confirmed to pertain to all workers (Table 6). The positive association between AI for decisions and HPWP is also qualified. Only mid- and high-skill workers experience the positive relationship between HPWP1 and AI for decisions, and only high-skill workers experience the positive relationship between HPWP2 and AI for decisions. Finally, AI orders is now found to decrease autonomy for the high-skill workers.

All in all, these results show that AI use has differing effects on work and job skill requirements depending on the type of use and the skill level of the job:

High-skill jobs that involve daily use of AI for decision making are very different from high-skill jobs that do not. Such jobs have less monotony and more learning and more use of HPWP, both in terms of teamwork and job rotation (HPWP1) and quality control and meeting quality standards (HPWP2). However, such jobs also have more work pace constraints.

High-skill jobs that involve receiving orders daily from an AI system, however, experience only increased work pace constraints and decreased autonomy.

Mid-skill jobs that involve daily use of AI for decision making differ from other mid-skill jobs that do not by having more use of HPWP1 (task rotation and teamwork).

Mid-skill jobs that involve receiving orders from AI daily have more constraints, more monotony and less learning than other mid-skill jobs.

All jobs, regardless of skill level, have more pace constraints when daily receiving orders from an AI system.

Therefore, the positive effect on skill requirements associated with AI for decision making (adaptability, social interaction, judgement) are primarily felt by high-skill workers. Such skill requirements arguably make the jobs more interesting and potentially also more productive and hence may allow for a higher wage. The negative effect of AI orders, however, are felt across jobs but are even stronger among mid-skill workers who also experience decreased learning and increased monotony. Even the general relationship that AI orders is associated with increased constraints will affect mid-skill jobs more, as AI for orders is used relatively often in such jobs; cf. Figure 1. The result that AI use affects workers in mid-skill jobs particularly adversely is consistent with AI use contributing to job polarization both in terms of wages and jobs, as mid-skill jobs are simplified.

Relatively few statistically significant relationships between AI use and skill requirements in low-skill jobs were found in the regressions. One possible explanation is the different nature of decisions made in jobs across skill levels. While decision making in a high-skill job may include which tasks to undertake, decision making in a low-skill job could relate more to the order of tasks and hence not affect which specific skills are needed.

4.3 The causal relationship between AI and skill requirements

The above results clearly show a relationship between the use of AI and skill requirements, but as the data are collected with a single survey and both questions on tasks at work and questions on technology use at work pertain to the respondent’s current main paid job, the results do not show the direction of causality—if any—between technology use and skill requirements. This is a general shortcoming of survey data. To use our results as a basis for managerial action, however, indications of the direction of causality are very useful. The identification of causality generally relies on the use of instrumental variables and untestable assumptions regarding the exogeneity of such instruments. The survey data used for our analysis do not provide us with a useable instrument. However, ANM, a tool used in ML to tease out causality between otherwise simultaneous variables, can also be applied for causal inference from cross-sectional survey data (Coad et al., 2018). We are therefore testing for causality in a narrow sense, as we are testing for consistency with causality as described by the ANM.

ANM identifies causality from X to Y if the relationship $Y = f(X)$ is consistent with an ANM while the reverse relationship, $X = f(Y)$, is not. $f(\cdot)$ is the ANM and is what in econometrics is a linear additive regression equation with an error term. $Y = f(X)$ is consistent with an ANM if the relationship does not exhibit endogeneity, which is identified by testing whether the residual is independent of X . The test for independence between X and the residual is the Hilbert Schmidt Independence Criterion (HSIC) test. The HSIC test is used to test for higher-order dependencies between the residual and X as opposed to just linear dependence. Whether the method supports or rejects causality is interpreted from a comparison of the p-values from modelling causality in each direction. If there is endogeneity in both or neither direction, then the method is inconclusive (Coad et al., 2018).

As the usefulness of ANM for discrete data is not well known, we do not use the binary indicators of daily use of AI as used in the regressions above. Instead, we use a variable that takes the values 1–5 for every day, at least once per week, 1–3 per month, less than once per month and never. There are 10 potential causal relations to test—the two variables for AI use by the five factors for tasks at work. For each test, we first correct the left-hand side variable for the variance that can be explained by the variables included as controls in the regressions—occupation, age, industry and education.

The p-values from the resulting 20 HSIC tests are shown in Table 7.

[Table 7 here]

In Table 7, the cells with dark shading provide relatively conclusive evidence, as the p-value in one direction is relatively large while the other p-value is small. Interpreted with the sign of the relationship as identified in the regressions, this means that AI for decision making causes lower monotony and more learning, while AI orders cause increased monotony and HPWP1 and lower autonomy. The lightly shaded cells are relationships where it depends on the p-value cut-off used for interpretation whether the result is conclusive. In all cases, there is strong evidence of endogeneity when modelling causality as an ANM from work tasks to AI, while the evidence is somewhat weaker in the opposite direction. It can therefore be argued that AI for decision making causes increased constraints, HPWP1 and HPWP2. The relationship between AI orders and HPWP1 seen in Table 7 was not observed in the regressions. It is possible that there is a non-linear relationship between these two variables, which we are not able to capture in the regressions.

5 Conclusions

The analysis in this paper using a novel and unique dataset on the use of AI among employees in Denmark allows us to engage with some of the limitations of the existing literature and the task-based approach. Unlike most of the literature investigating the impact of AI on skills, the analysis presented here does not assume that the effects of AI on skills are uniform for all employees with the same occupational title. Rather, our micro-level data allows us to handle each job separately and investigate within-job relationships between technology use and skill requirements as revealed with the JRA approach. This allows us to go beyond the existing literature, which has assumed for the most part that the effect of AI is to substitute for existing skills, and to show that the effects of AI are varied and depend on how it is used. AI may enhance or augment skills in the sense of increasing complexity and learning, as well as the need for social interactive and adaptability skills, or it may simply increase work pace constraints and reduce employee autonomy. We use ANM to establish the likely direction of causality in our results and find that the direction of causality is from AI use to skill requirements. This result is not found consistently in all cases, but it is contradicted in none of our cases.

In our data, we capture two different types of use of AI from the perspective of the employee without distinguishing between symbolically coded AI and AI created from ML. One type is where the AI has chosen the action to be performed and gives the employee orders or instructions, while the other type is where the employee uses information from the AI in his or her own decision making. These two types of AI use are relatively generic, which they must necessarily be when conducting a broad survey of AI use. Therefore, our analysis does not approach issues of domain-specific skills.

For example, in the case of high-skilled jobs our regression analysis shows that while the use of AI as an input for further decision making entails an increase in job pace constraints it also results in less monotony and more learning and in increased use of a wide range of HPWP, including teamwork, job flexibility and responsibility for quality control. However, if AI is used as a tool for giving orders to workers with high-skill jobs it only results in increased work pace constraints and decreased autonomy. Similar differences in AI impacts on skills are found to be true for medium-skill jobs, although to a lesser degree than in the case of high-skill jobs. For the workers with medium-skill jobs, fewer of the HPWP practices are associated with using AI for decision making, and the adverse effects of using AI as a tool for giving orders are enhanced in the sense of a larger positive impact on work pace constraints, more monotony and less learning. For low-skill jobs, the effects of using AI as an input for further decision making are neutral, while the dominant effect of using AI for giving orders is to increase work pace constraint.

Our results clearly are not consistent with the literature arguing that the adoption of AI foretells an economic singularity that will uniformly reduce the importance of human skills in work activity. The observed variability in how AI is used and in its impacts raises the important question of how much discretion or voluntarism employers have in how they use AI and whether they can decide to use AI in ways that enhance skills and assure that work is less monotonous and provides more opportunities for learning. The answer to this question in turn has important implications for the policies employers adopt for their provision of AI-related training and skills development on and off the job. It also has important implications for the scope that the social partners have for shaping labour market policies that serve to 'future-proof' skills and competence development by providing opportunities for workers to continuously renew and upgrade their skills. The Danish institutional framework includes unions that traditionally have been focussed on maintaining jobs through supporting efforts to improve firms' competitiveness and hence partake in competence upgrading both through workers' rights to further

education and training explicitly defined in collective agreements, and corporatist policy making aimed at developing the active labour market policy (Ibsen, 2012). The distinctive Danish system of flexible security which combines unemployment protection with public provision of training for the unemployed, also supports skills and competence development. As discussed in Nielsen et al. 2021, this objective can be furthered by actively involving local educational institutions, and one possible solution suggested by the authors is establishing local competence clusters where employers and unions in partnership with educational institutions are systematically involved in training and re-training of the workforce.

We see a pressing need for additional micro-level data to better address these issues. Employee-level survey data complementary to that collected in the TASK survey could provide information on the types and amount of training received by employees. Assessing how much discretion employers have in how AI is used and in the role of worker representatives in monitoring skills needs and in directing the demands made on external training providers could be explored with complementary micro-data at the enterprise level. It is our hope that the analysis presented here will encourage both researchers and policy makers alike to pursue this data collection effort and to contribute to generating information that can improve our understanding of AI and how it may be mobilized as a tool to both increase the quality of working life and economic performance.

Data availability statement

The data that support the findings of this study are available from the corresponding author upon reasonable request and subject to conditions and approval by Statistics Denmark.

References

- Acemoglu D and Restrepo P (2017) Robots and jobs: Evidence from US labor markets, National Bureau of Economic Research, Working Paper 23285. Cambridge MA.
- Aghion P, Jones BF and Jones CI (2019) 'Artificial Intelligence and Economic Growth' in *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, Chicago, pp. 237-282.
- Agrawal A, Gans J and Goldfarb A (Eds.) (2019) *The Economics Of Artificial Intelligence: An Agenda*. University of Chicago Press, Chicago.
- Appelbaum E, Bailey T, Berg P and Kalleberg AL (2000) *Manufacturing Advantage: Why High-Performance Work Systems Pay Off*. Cornell University Press, Ithaca, NY.
- Arntz M, Gregory T and Zierahn U (2016) The risk of automation for jobs in OECD countries: A comparative analysis. *OECD Social, Employment and Migration Working Papers*, No. 189. OECD Publishing, Paris.
- Athey S (2019) 'The Impact of Machine Learning on Economics', in *The Economics Of Artificial Intelligence: An Agenda*. University of Chicago Press, Chicago, pp. 507-547.
- Autor DH, Levy F and Murnane RJ (2003) The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics* 118, 4: 1279-1333.
- Autor DH and Price B (2013) The changing task composition of the US labor market: An update of Autor, Levy, and Murnane (2003), mimeo. Available from: <https://www.brendanmichaelprice.com/research/other/Autor-Price-2013.pdf>.
- Autor D (2015) Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3): 3-30.
- Black SE and Lynch LM (2004) What's driving the new economy?: The benefits of workplace innovation. *The Economic Journal* 114, 493, F97-F116.
- Braverman H (1974) *Labor and Monopoly Capital: The Degradation of Work in The Twentieth Century*. NYU Press, New York.
- Brynjolfsson E and McAfee A (2014) *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company, New York.
- Brynjolfsson E, Mitchell T and Rock D (2018) What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings* 108: 43-47.
- Caroli, E., & Van Reenen, J. (2001). Skill-biased organizational change? Evidence from a panel of British and French establishments. *The Quarterly Journal of Economics*, 116(4), 1449-1492.

- Coad A, Janzing D and Nightingale P (2018) Tools for causal inference from cross-sectional innovation surveys with continuous or discrete variables: Theory and applications. *Cuadernos de Economía* 37(SPE75): 779-807.
- Dauth W, Findeisen S, Südekum J and Woessner N (2017) German robots - The impact of industrial robots on workers, IAB Discussion Paper 30/2017. Institut für Arbeitsmarkt und Berufsforschung (IAB), Nürnberg.
- Felstead A, Gallie D and Green F (2002) *Work Skills in Britain, 1986-2001*. Department for Education and Skills, London.
- Edwards P and Ramirez P (2016) When should workers embrace or resist new technology? *New Technology, Work and Employment*, 31(2): 99-113.
- Felten EW, Raj M and Seamans R (2019) The occupational impact of artificial intelligence: Labor, skills, and polarisation. NYU Stern School of Business. Available from: <https://ssrn.com/abstract=3368605>.
- Freeman RB and Kleiner MM (2000) Who benefits most from employee involvement: Firms or workers? *American Economic Review* 90(2): 219-223.
- Frey CB and Osborne MA (2013) The future of employment: How susceptible are jobs to computerisation? Working Paper. Oxford Martin Programme on Technology and Employment, Oxford.
- Frey CB and Osborne MA (2017) The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change* 114: 254-280.
- Gjerding AN, Holm JR, Lorenz E and Stamhus J (2020) Ready, but challenged. *Journal of Business Models working paper*, 1, 001.
- Goos M and Manning A (2007) Lousy and lovely jobs: The rising polarisation of work in Britain. *The Review of Economics and Statistics* 89(1): 118-133.
- Goos M, Manning A and Salomons A (2009) Job polarization in Europe. *American Economic Review* 99(2): 58-63.
- Graetz G and Michaels G (2018) Robots at work. *Review of Economics and Statistics* 100(5): 753-768.
- Guest DE (2017). Human resource management and employee well-being: Towards a new analytic framework. *Human Resource Management Journal* 27(1): 22-38.
- Giuntella O and Wang T (2019) Is an army of robots marching on Chinese jobs? IZA Discussion Papers, No. 12281. IZA, Bonn.

- Holm JR and Lorenz E (2015) Has “discretionary learning” declined during the Lisbon Agenda? A cross-sectional and longitudinal study of work organization in European nations. *Industrial and Corporate Change* 24(6): 1179-1214.
- Holm JR, Lorenz E and Nielsen P (2020) Work organization and job polarization. *Research Policy* 49(8): 104015.
- Holm JR, Lorenz E and Stamhus J (2021) ‘The Impact of Robots and AI/ML on Skills and Work Organization’ in Christensen JL, Gregersen B, Holm JR and Lorenz E *Globalization, New and Emerging Technologies and Sustainable Development – The Danish Innovation System in Transition*. Routledge, London.
- Ibsen, CL (2012). *Trade unions in Denmark*. Berlin: Friedrich-Ebert-Stiftung.
- Katz, LF, & Murphy, KM (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1), 35-78.
- Katz LF and Autor DH (1999) ‘Changes in the Wage Structure and Earnings Inequality’ in *Handbook of Labor Economics*, vol. 3. Elsevier, Amsterdam, pp. 1463-1555.
- Keynes JM (1930) ‘Economic possibilities for our grandchildren’ in Keynes, J.M., *Essays in persuasion* (2nd ed.). Macmillan, London.
- Lloyd C and Payne J (2019) Rethinking country effects: Robotics, AI and work futures in Norway and the UK. *New Technology, Work and Employment* 34(3): 208-225.
- Lorenz E and Potter J (2019) Workplace organisation and innovation in small and medium-sized enterprises. *OECD SME and Entrepreneurship Papers*, No. 17. OECD Publishing, Paris.
- Machin, S., & Van Reenen, J. (1998). Technology and changes in skill structure: evidence from seven OECD countries. *The Quarterly Journal of Economics*, 113(4), 1215-1244.
- Nielsen P, Holm JR and Lorenz E (2021) ‘Work Policy and Automation in the Fourth Industrial Revolution’ in Christensen JL, Gregersen B, Holm JR and Lorenz E *Globalisation, New and Emerging Technologies, and Sustainable Development – The Danish Innovation System in Transition*. Routledge, London.
- Noon M, Blyton P and Morrell K (2013) *The Realities of Work: Experiencing Work and Employment in Contemporary Society*. Macmillan International Higher Education.
- Nordhaus WD (2015) Are we approaching an economic singularity? Information technology and the future of economic growth’, No. w21547. National Bureau of Economic Research.
- OECD (2009) PIAAC Background Questionnaire, JRA version 5.0 Conceptual Approach. OECD Publishing, Paris.
- Piva, M, Santarelli, E, & Vivarelli, M (2005). The skill bias effect of technological and organisational change: Evidence and policy implications. *Research Policy*, 34(2), 141-157.

Sanders, M, & Ter Weel, B (2000). Skill-biased technological change: theoretical concepts, empirical problems and a survey of the evidence. *DRUID Papers*, 8.

Group	Label	Percent of employees	Skill level	Job examples
1	Managers	5.59	High	CEO, vice president, sales/factory/school manager
2	Professionals	34.36	High	Professor, engineer, doctor, nurse, teacher, chief consultant, computer programmer, auditor, lawyer
3	Technicians and associate professionals	16.99	High	Consultant, lab assistant, pharmacist, accountant, technician, head of department
4	Clerical support	10.99	Mid	Office assistant, administrative officer, secretary, receptionist, payroll clerk
5	Service and sales	11.26	Low	Kitchen assistant, shop assistant, building caretaker, personal care assistant, porter, watchman
7	Craft and related trades	6.70	Mid	Brick layer, carpenter, smith, electrician
8	Plant and machine operators and assemblers	6.81	Mid	Chauffeur, machine operator in industry/meat/boiler/cutting, locomotive driver
9	Elementary occupations	7.31	Low	Cleaning, gravedigger, tradesman's assistant, warehouse assistant

Note: Data and examples from the TASK survey. Post-stratification weights are used.

Table 1: Occupational groups

Variable	Percent of employees
Learning new things in work	60.4
Complexity of tasks	64.8
Discretion in fixing work methods	56.8
Discretion in fixing work pace	57.1
Teamwork	58.3
Job rotation	57.0
Responsibility for quality control	86.0
Precise quality standards	80.4
Horizontal constraints on work pace	57.0
Hierarchical constraints on work pace	12.9
Norm-based constraints on work pace	42.1
Automatic constraints on work pace	15.4
Repetitiveness of tasks	33.3
Monotony	26.8

Table 2: Share of employees characterized by each work activity trait

Variable	Monotony	Constraints	HPWP1	Autonomy	HPWP2
Learning new things in work	-0.4362	0.1127	0.2475	0.0253	0.005
Complexity	-0.4471	0.1323	0.1168	-0.0013	0.1854
Discretion in fixing work methods	-0.1244	0.0857	0.0964	0.5813	-0.0573
Discretion in fixing work pace	0.0765	-0.0103	-0.0279	0.7101	0.0806
Teamwork	0.0823	-0.0783	0.5410	0.1877	0.0650
Job rotation	-0.0229	-0.0823	0.6130	-0.0732	-0.0987
Responsibility for quality control	-0.0361	-0.0500	-0.0650	0.0830	0.7826
Precise quality standards	0.0890	0.1484	0.1428	-0.2814	0.4889
Horizontal constraints on work pace	-0.0505	-0.2904	0.3293	-0.1481	0.1926
Hierarchical constraints on work	-0.0481	0.5673	-0.1617	0.0196	-0.1591
Norm-based constraints on work pace	-0.0478	0.5139	0.0790	0.0060	0.1412
Automatic constraints on work pace	0.1894	0.4877	-0.0890	0.0874	0.0350
Repetitiveness of tasks	0.4925	0.0338	0.2659	0.0238	0.0815
Monotony	0.5352	0.1248	0.0509	0.0080	-0.0158

Table 3: Factor loadings

Table 4: Regression results

	Model 1 + Monotony - Learning	Model 2 Constraints	Model 3 HPWP1	Model 4 Autonomy	Model 5 HPWP2
AI.decisions	-0.151 (0.13)	0.263* (0.14)	0.531*** (0.13)	0.005 (0.12)	0.300** (0.13)
AI.orders	0.122 (0.16)	0.867*** (0.18)	0.020 (0.17)	-0.117 (0.14)	0.200 (0.15)
Occupation high		Reference			
Occupation middling	0.717*** (0.13)	-0.067 (0.13)	-0.384*** (0.13)	-0.166 (0.11)	-0.129 (0.11)
Occupation low	0.476*** (0.12)	-0.276** (0.13)	-0.285** (0.12)	-0.231** (0.12)	-0.210* (0.12)
Education primary		Reference			
Education secondary or <=2 years tertiary	-0.167 (0.13)	-0.337** (0.17)	0.005 (0.14)	0.038 (0.14)	0.080 (0.12)
Education >2 years tertiary	-0.752*** (0.16)	-0.448** (0.19)	0.171 (0.16)	0.104 (0.15)	0.025 (0.15)
Age 18-39		Reference			
Age 40-59	-0.296*** (0.10)	-0.294*** (0.10)	-0.130 (0.10)	0.113 (0.09)	0.146* (0.09)
Age => 60	-0.427*** (0.12)	-0.465*** (0.14)	-0.448*** (0.14)	0.150 (0.13)	0.310*** (0.10)
Agriculture, Forestry and Fishing		Reference			
Industry and Mining	-0.638** (0.32)	0.104 (0.47)	-0.217 (0.46)	-0.184 (0.31)	-0.198 (0.26)
Construction	-1.223*** (0.36)	-0.263 (0.50)	0.043 (0.48)	-0.133 (0.35)	0.098 (0.28)
Trade and Transport	-0.398 (0.31)	-0.191 (0.47)	-0.338 (0.46)	-0.486 (0.31)	-0.445* (0.26)
Information and Communication	-0.882** (0.35)	-0.324 (0.50)	-0.032 (0.50)	-0.399 (0.34)	-0.401 (0.31)
Finance, Insurance	-0.824**	-0.453	-0.531	-0.699**	-0.722**

and Real Estate	(0.37)	(0.48)	(0.48)	(0.34)	(0.30)
Business services	-0.545*	-0.291	-0.327	-0.495	-0.058
	(0.32)	(0.47)	(0.47)	(0.32)	(0.25)
Education and Health	-0.587*	-0.503	-0.024	-0.670**	-0.419*
	(0.31)	(0.46)	(0.46)	(0.30)	(0.25)
Culture, Leisure and other services	-0.560	-0.472	-0.264	-0.160	-0.513
	(0.35)	(0.52)	(0.50)	(0.37)	(0.33)
Constant	1.016***	0.819*	0.269	0.442	0.152
	(0.35)	(0.50)	(0.48)	(0.34)	(0.29)
<hr/>					
N	1,116	1,116	1,116	1,116	1,116
R-squared	0.192	0.113	0.077	0.035	0.056

* p<0.10, ** p<0.05, *** p<0.01; Robust standard errors in parentheses. The data are weighted.

	Monotony	Constraints	HPWP1	Autonomy	HPWP2
High	-0.289* (0.159)	0.390** (0.177)	0.340*** (0.148)	0.029 (0.149)	0.327** (0.165)
Middling	-0.012 (0.276)	0.141 (0.294)	0.930*** (0.220)	0.140 (0.236)	0.390 (0.263)
Low	-0.210 (0.350)	-0.112 (0.275)	0.486 (0.408)	-0.218 (0.260)	0.130 (0.344)

Robust SE in parentheses. Full results in appendix.

Table 5: Marginal effects of AI for decisions by occupational group

	Monotony	Constraints	HPWP1	Autonomy	HPWP2
High	0.011 (0.249)	0.679*** (0.254)	-0.204 (0.215)	-0.528** (0.225)	0.258 (0.200)
Middling	0.607** (0.292)	1.238*** (0.363)	-0.197 (0.335)	0.039 (0.266)	-0.018 (0.298)
Low	-0.326 (0.278)	0.776** (0.307)	0.424 (0.298)	0.241 (0.221)	0.332 (0.301)

Robust SE in parentheses. Full results in appendix.

Table 6: Marginal effects of AI orders by occupational group

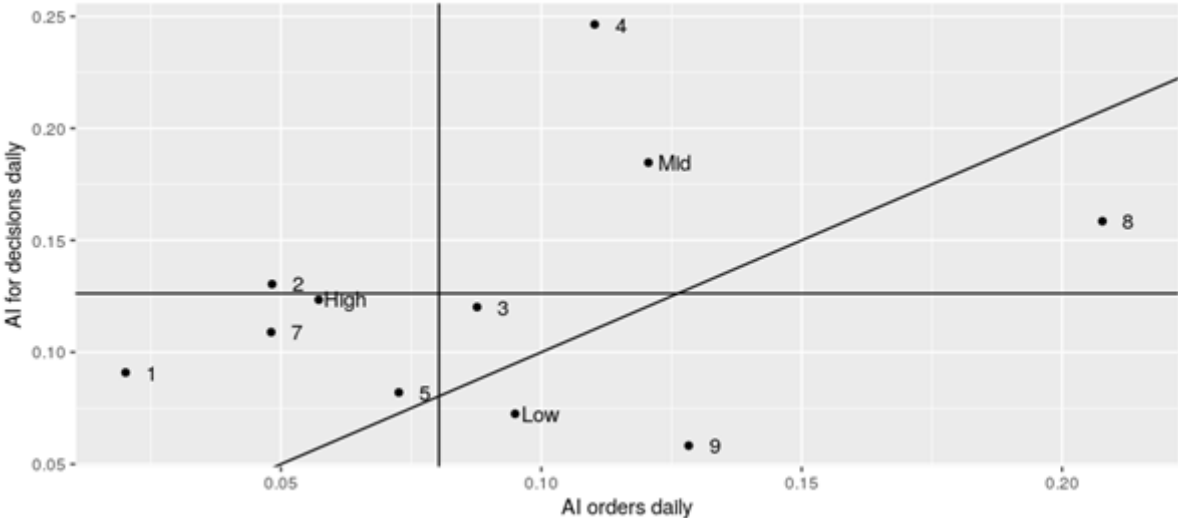
Factor	Monotony	Constraints	HPWP1	Autonomy	HPWP2
From AI orders	0.119	<0.001	0.960	0.450	<0.001
To AI orders	0.005	<0.001	0.002	0.014	<0.001
From AI for decision	0.456	0.070	0.068	0.173	0.020
To AI for decision	0.029	<0.001	<0.001	0.129	<0.001

p-values from ANM HSIC test. 'From'/'To' is the implied direction of causality

Table 7: Results from additive noise modelling

Figure captions

Figure 1: Daily AI use by occupational group and type of use



Appendix

Below are the 14 questions from the TASK survey that are used to construct the factors for skills used at work. All 14 questions are answered using the scale Always-Often-Sometimes-Rarely-Never. We use only binary indicators coded 1 if the response was 'Always' or 'Often'.

How often does your main job involve:

1. Meeting precise quality standards?
2. Assessing yourself the quality of your own work?
3. Monotonous tasks?
4. Complex tasks?
5. Learning new things?
6. Short, routine and repeated tasks of less than 10 minutes?
7. Rotating tasks between yourself and colleagues?
8. That you are able to choose or change your methods of work?
9. That you have the option to change your speed of work?
10. That you work in a group or team that has common tasks and plan its work?

How often does your pace of work depend on:

11. The work done by colleagues?
12. Measurable production targets or performance targets?
13. Automatic speed of a machine or movement of a product?
14. Direct control of superiors?

Regression with interaction effects

Table 8 shows the results of expanding the regression in equation 1 to include interaction terms between the variables for AI use and the categorical variable for occupation. This entails that the marginal effects of AI use varies across occupational groups. The estimate for the direct effect on AI use then becomes the estimate for the reference group High, as no interaction terms are included when computing this marginal effect. The estimates for the interaction terms are the differences in the marginal effects between occupational groups, and this must be added to the direct effect (i.e. the effect for High) when computing the marginal effects for other groups. This can also be seen from taking the partial derivative of the regression equation. For example, the marginal effect of AI.orders on Monotony is

$$\frac{\partial Monotony_i}{\partial AI.orders_i} = \alpha_{AI.orders} + \beta_{AI.orders*Midling} Midling_i + \beta_{AI.orders*Low} Low_i. \quad (2)$$

In equation 2, Monotony is the first factor from the PCA and AI.orders, Midling and Low are all binary variables as described in the main text.

[Table 8 here]

Table 9 below shows the value of the eigenvalues for the five components after orthogonal varimax rotation.

[Table 9 here]

Table 8: Regression results

	Model 1 + Monotony - Learning	Model 2 Constraints	Model 3 HPWP1	Model 4 Autonomy	Model 5 HPWP2
AI.decisions	-0.289* (0.16)	0.390** (0.18)	0.400*** (0.15)	0.029 (0.15)	0.327** (0.16)
AI.orders	0.011 (0.25)	0.679*** (0.25)	-0.204 (0.21)	-0.528** (0.23)	0.258 (0.20)
Occupation high	Reference				
Occupation middling	0.607*** (0.14)	-0.085 (0.14)	-0.468*** (0.14)	-0.236* (0.13)	-0.113 (0.12)
Occupation low	0.499*** (0.12)	-0.236* (0.14)	-0.351*** (0.12)	-0.272** (0.13)	-0.204 (0.13)
AI.decisions * middling	0.276 (0.32)	-0.249 (0.34)	0.530** (0.26)	0.111 (0.28)	0.063 (0.31)
AI.decisions * low	0.079 (0.38)	-0.502 (0.33)	0.086 (0.44)	-0.247 (0.30)	-0.197 (0.38)
AI.orders * middling	0.596 (0.38)	0.560 (0.44)	0.006 (0.40)	0.568 (0.35)	-0.276 (0.36)
AI.orders * low	-0.336 (0.37)	0.097 (0.39)	0.628* (0.37)	0.769** (0.31)	0.074 (0.36)
Education primary	Reference				
Education secondary or <=2 years tertiary	-0.172 (0.13)	-0.341** (0.17)	0.011 (0.14)	0.042 (0.14)	0.080 (0.12)
Education >2 years Tertiary	-0.756*** (0.16)	-0.452** (0.19)	0.168 (0.16)	0.099 (0.16)	0.024 (0.15)
Age 18-39	Reference				
Age 40-59	-0.295*** (0.10)	-0.296*** (0.10)	-0.127 (0.10)	0.112 (0.09)	0.149* (0.09)
Age => 60	-0.460*** (0.12)	-0.484*** (0.13)	-0.447*** (0.14)	0.140 (0.13)	0.317*** (0.10)

Agriculture, Forestry and Fishing			Reference		
Industry and Mining	-0.671** (0.32)	0.109 (0.46)	-0.222 (0.47)	-0.188 (0.31)	-0.179 (0.26)
Construction	-1.199*** (0.36)	-0.246 (0.50)	0.084 (0.50)	-0.092 (0.35)	0.105 (0.28)
Trade and Transport	-0.436 (0.32)	-0.193 (0.47)	-0.328 (0.48)	-0.476 (0.31)	-0.432 (0.27)
Information and Communication	-0.887** (0.35)	-0.341 (0.50)	-0.028 (0.51)	-0.416 (0.34)	-0.393 (0.31)
Finance, Insurance and Real Estate	-0.809** (0.37)	-0.449 (0.47)	-0.507 (0.50)	-0.677** (0.34)	-0.724** (0.31)
Business services	-0.583* (0.33)	-0.282 (0.46)	-0.319 (0.48)	-0.477 (0.32)	-0.043 (0.26)
Education and Health	-0.619** (0.31)	-0.497 (0.45)	-0.020 (0.47)	-0.659** (0.30)	-0.407 (0.26)
Culture, Leisure and other services	-0.550 (0.36)	-0.428 (0.52)	-0.241 (0.51)	-0.109 (0.37)	-0.501 (0.34)
Constant	1.072*** (0.36)	0.815 (0.50)	0.289 (0.50)	0.455 (0.34)	0.133 (0.29)
N	1,116	1,116	1,116	1,116	1,116
R-squared	0.201	0.116	0.084	0.042	0.057

* p<0.10, ** p<0.05, *** p<0.01; Robust standard errors in parentheses. The data are weighted

Table 9: Eigenvalues of first 5 rotated components

Component	Variance	Proportion	Cumulative
Component 1	1.70211	0.1216	0.1216
Component 2	1.66715	0.1191	0.2407
Component 3	1.62262	0.1159	0.3566
Component 4	1.44189	0.1030	0.4596
Component 5	1.13361	0.0810	0.5405