Automation and the Changing Composition of Skill Demand*

Mark Hellsten[1], Giuseppe Pulito[†2], and Sarah Schroeder[3,4]

[1]University of Tübingen
[2]ROCKWOOL Foundation Berlin
[3]Aarhus University
[4]Ratio Institute

June 12, 2026

Abstract

This paper provides new evidence on how automation reshapes firms' demand for skills, not only by changing the occupational composition, but also by reshaping what existing jobs require. Using matched data on firm-level automation investments and detailed job vacancy postings from Denmark, we extract multidimensional skill profiles through natural language processing and decompose changes in skill demand into within- and between-occupation components. Within-occupation adjustment is a quantitatively important margin, accounting for 14–39% of total skill demand change depending on skill type and occupational group. Drawing on a task-based framework that links automation to shifts in multiple skill types within occupations, we estimate the causal effect of automation using a staggered difference-in-differences design. The effects are heterogeneous across the occupational hierarchy: among managers and professionals, automation increases the demand for soft skills, shifting the within-occupation skill mix toward interpersonal and cognitive competencies; among production workers, adjustment operates primarily through reduced hiring rather than changes in skill requirements, while retraining intensity rises by 5 percentage points. Our findings highlight that automation operates through multiple adjustment margins, with implications for training policy and labour market resilience.

Keywords: automation, skills, task content, labour demand, technological change, job vacancies, within-occupation adjustment

JEL classification: J24, O33, M51, L23

*We thank an editor and three anonymous referees for helpful comments. We are grateful to Rosario Crinò, Angelo Cuzzola, Georg Graetz, Anders Humlum, Ina Jäkel, Michael Koch, Magnus Lodefalk, Gabriele Lucchetti, and Lars Skipper for valuable comments and discussions. We are also grateful for feedback from seminar participants at AI-Econ Lab, the Ratio Institute, workshop participants at Workshop on the Impact of Automation on Labor Markets 2025, Aarhus-Kiel Workshop 2024, DGPE workshop 2024, and Internal FIND Workshop 2024. Pulito acknowledges financial support from the ROCKWOOL Foundation Berlin (Project No. 2002). Financial support from the Carlsberg Foundation, the European Union through the European Research Council (ERC-2022-SG ORGANDICT, No. 101076790) and the Ratio Institute is gratefully acknowledged.

†Corresponding author: Giuseppe Pulito, Email: gp@rfberlin.com

1 Introduction

Automation technologies are transforming the nature of work across industrialized economies. As firms adopt new machinery and digital tools, production processes are reorganized and tasks are reallocated between workers and machines (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2020). These changes reshape the skill content of jobs, with consequences for wage inequality, worker displacement, and the returns to different types of human capital (Autor, 2015; Acemoglu and Restrepo, 2022). A large body of research has studied how technological change affects employment across occupations, documenting patterns of job polarization and skill-biased demand shifts (Autor and Dorn, 2013; Goos et al., 2014). Yet considerably less is known about how automation changes what existing jobs require, specifically, whether firms adjust primarily by reallocating employment across broad occupational groups, or by reconfiguring the bundle of skills demanded *within* occupations.

Distinguishing between these two margins of adjustment is central to understanding how technological change affects workers and how firms adapt. If automation operates primarily through occupational reallocation, standard measures of employment shifts across job categories capture the relevant dynamics. But if the skill content of existing jobs changes substantially in response to automation, with some skills gaining importance while others become obsolete, then policies focused on occupational mobility may miss a key channel through which technology reshapes work. Moreover, the nature of within-occupation task change may vary across the occupational hierarchy: occupations differ in their task composition, and the same automation event may displace routine activities in some roles while raising the importance of complementary, non-routine tasks in others. Understanding these potentially heterogeneous within-occupation adjustments is essential for designing effective training and reskilling policies.

In this paper, we study how firms adjust the skill content of jobs following automation

events, distinguishing between within-occupation changes in skill demand and reallocation across occupations. Using rich Danish data that link firm-level automation events to detailed job vacancy postings and administrative worker-level records, we show that automation leads to sharply different adjustment patterns across the occupational hierarchy, and that within-occupation changes in the composition of skill demand, rather than occupational restructuring, are a quantitatively important margin of adjustment.

To interpret the empirical patterns, we develop a task-based conceptual framework, inspired by canonical models of automation (Acemoglu and Restrepo, 2020), that incorporates multidimensional skill demand within occupations. Standard task-based approaches model how automation reallocates tasks between labour and capital, but typically characterize skill demand along a single dimension or at a high level of aggregation. Our framework extends this task-based approach by defining occupation-level skill shares as aggregates over a latent task distribution, allowing automation to shift the composition of multiple skill types within occupations through two channels: displacement of tasks intensive in automatable skills, and reweighting of surviving tasks toward complementary activities. This yields predictions about within-occupation changes in the demand for different types of skills, which we take to the data.

Our empirical strategy leverages detailed firm-level records of machinery investments to identify discrete automation events, following the investment-spike approach of Bessen et al. (2023) and Békés et al. (2025). We combine these events with an extensive database of job postings from Denmark's largest online vacancy platform, from which we extract multidimensional skill requirements for soft skills, high-complex hard skills, and low-complex hard skills, building on the skill classification frameworks of Deming (2017) and Deming and Noray (2020). Following the conceptual framework, we decompose skill changes into between-occupation reallocation and within-occupation change. We estimate the causal effect of automation on the composition of skill demand using the staggered difference-in-differences estimator of Callaway and Sant'Anna (2021), restricting comparisons to not-yet-treated firms

to address concerns about systematic differences between automating and non-automating firms. To assess whether these changes translate into actual workforce adjustment, we complement the vacancy analysis with linked register data on worker experience, education, and participation in retraining courses.

Results from a descriptive decomposition of skill demand around automation events reveal substantial heterogeneity across occupations in the direction and magnitude of within-occupation adjustment. The within-occupation component accounts for a sizable share of total skill demand change, ranging from 14% to 39% depending on skill type and occupational group, underscoring that automation operates not only through occupational reallocation but also through changes in skill demand within occupations. Our estimates confirm that automation has heterogeneous effects on skill demand across occupations and skill types. Among managers & professionals, it increases demand for soft skills by 8 percentage points, shifting the within-occupation skill mix toward interpersonal and cognitive competencies. By contrast, for production & elementary workers, we find no statistically significant within-occupation changes in vacancy skill demand; instead, adjustment operates primarily through lower hiring and the resulting contraction in employment.

Evidence from complementary register data on workforce characteristics suggests that changes in skill demand are accompanied by adjustments in firms' actual workforce. We find that the increased demand for soft skills is weakly corroborated by a corresponding increase in experience among managers & professionals, consistent with firms relying more on tasks requiring judgment and coordination, typically accumulated through labour market experience rather than education. Within production & elementary occupations, we observe a rise in participation in adult retraining courses by 5 percentage points, suggesting that firms invest in targeted reskilling to equip incumbent workers with the competencies required alongside new machinery.

Together, these findings show that within-occupation changes in skill demand are an impor-

tant channel through which firms adapt to automation, complementing occupational reallocation. They also reveal an asymmetry across the occupational hierarchy: higher-tier roles increasingly depend on interpersonal skills and experience, while lower-tier jobs rely more on targeted retraining than on new skills demanded in the labour market. This suggests that policies narrowly focused on occupational mobility or broad tertiary education expansion may overlook how firms adjust in practice. Instead, our results highlight the importance of continuous, occupation-specific reskilling, an approach well aligned with Denmark's emphasis on lifelong learning, to help workers adapt within their current roles and mitigate the uneven burden of technological change.

Our study contributes to four strands of the literature. First, it relates to research examining how technologies alter the task content of work (Autor et al., 2003; Spitz-Oener, 2006; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018, 2020). This literature shows that technology reshapes labour demand by displacing routine, codifiable tasks and complementing non-routine activities, though the focus has primarily been on changes in employment across occupations or broad skill groups rather than on how these forces alter the composition of skills within occupations. We contribute by developing and empirically documenting a task-based framework that links task reallocation to multidimensional skill demand within occupations, allowing us to distinguish between displacement and complementarity mechanisms in observed skill changes.

Second, our paper contributes to a growing literature on how automation alters occupational expertise requirements. Autor and Thompson (2025) show that task removal can raise or lower expertise needs, with opposing effects on wages and employment. Salomons et al. (2025) find that technology exposure leads to within-occupation changes in skill supply, while Caines et al. (2017) highlight that technological change favors complex tasks. We extend this work by combining firm-level automation events with vacancy and register data, and show that automation shifts skill demand toward interpersonal and cognitive competencies within high-tier roles and induces retraining in lower-tier jobs. The increased reliance on

experienced workers among managers & professionals is consistent with automation raising the expertise threshold of these occupations, as tasks requiring less judgment are automated and the remaining activities demand greater accumulated competence.

Third, our research contributes to the literature on multidimensional skill dynamics, which argues that education or cognitive ability alone does not fully capture the key dimensions of worker skills (Lindqvist and Vestman, 2011; Heckman and Kautz, 2012; Grönqvist and Lindqvist, 2016; Deming, 2017; Deming and Kahn, 2017; Edin et al., 2022; Izadi and Tuhkuri, 2024; Deming and Silliman, 2025). These papers document the rising importance of cognitive and soft skills in the labour market. We contribute empirical evidence identifying automation as a key driver of changes in skill demand. Existing research on technology and skills has predominantly focused on computerization and AI (Autor et al., 1998; Acemoglu et al., 2022), with limited attention to the multidimensional nature of skills. Our findings show that automation affects the skill mix not only through occupational reallocation but also by altering skill demand within occupations. Related work on adjustment across multiple skill dimensions has instead focused on responses to recessionary shocks (Hershbein and Kahn, 2018).

Finally, our study relates to a broad literature on the labour market effects of technology adoption (Beaudry et al., 2016; Graetz and Michaels, 2018; Humlum, 2022; Aghion et al., 2022; Hirvonen et al., 2025; Bessen et al., 2023; de Souza and Li, 2023; Bonfiglioli et al., 2024). This literature documents heterogeneous employment effects of automation, including evidence of job polarization (Goos et al., 2014). Our findings highlight that these employment effects capture only part of the adjustment: automation also reshapes the content of work within jobs, altering the mix of skills demanded and inducing targeted workforce adjustments such as retraining. This underscores that technological change operates through multiple margins, namely employment, task composition, and skill demand, which need to be considered jointly.

The remainder of the paper is organized as follows. Section 2 presents our conceptual framework. Section 3 describes the data sources and outlines the construction of key variables, as well as provides a descriptive decomposition of skill demand in Subsection 3.5. Section 4 details the empirical strategy used to analyze the effects of automation on skill demand, and discusses results. Finally, Section 5 concludes.

2 Conceptual Framework: Automation, Task Reallocation and Multidimensional Skill Demand

This section develops a conceptual framework that links task-level automation decisions to shifts in occupation-level skill demand, both within and across occupations. Because our empirical analysis tracks changes in occupation-level skill requirements rather than task-level outcomes, the framework provides a mapping from task-level technological change to the occupation-level skill shares observed in the data.

A central premise is that skills are multidimensional and that automation does not affect all skill types symmetrically. A large body of evidence documents that worker productivity and labour market outcomes depend on multiple dimensions of human capital beyond formal education or cognitive ability alone (Heckman, 1995; Heckman and Kautz, 2012; Deming, 2017). At the same time, the literature on technological change has shown that automation most readily substitutes for routine, codifiable tasks, while tending to complement activities requiring judgment, coordination, and interpersonal interaction (Autor et al., 2003; Autor, 2015). These two observations jointly motivate our approach: because different types of tasks draw on different types of skills, automation-induced changes in task composition should translate into shifts in the relative demand for distinct skill types within occupations. We capture this by modeling skill demand in a multidimensional way. In the empirical analysis, we classify 15 individual skill categories extracted from job vacancy texts into three broad clusters: soft skills, high-complex hard skills, and low-complex hard skills.

Building on the task-based framework of Acemoglu and Restrepo (2020), we model how automation and task reallocation shift the composition of skill demand within and across occupations. Relative to standard task-based models, which typically characterize skill demand along a single dimension, the framework explicitly incorporates multiple skill types at the occupational level and generates predictions about how automation alters their relative importance.

2.1 Environment

We consider a continuum of tasks $i \in [0, 1]$.[1] Each occupation o is characterized by a distribution of task weights $\gamma_o(i)$, such that $\int_0^1 \gamma_o(i)\,\mathrm{d}i = 1$. The function $\gamma_o(i)$ captures the share of labour input (or time) allocated to task i within occupation o. Each task requires a vector of binary skill indicators $\mathbf{s}(i) = (s_1(i), \ldots, s_K(i))$, where $s_k(i) = 1$ if skill k is needed for task i. Each task also has an automatability index $a(i) \in [0, 1]$, with higher values indicating greater susceptibility to automation.

Firms perform tasks using either labour or capital. The cost of using labour to perform task i is:

$$\text{Cost}_L(i) = \sum_{k=1}^{K} s_k(i) v_k, \tag{1}$$

where v_k is the wage or price of skill k. The cost of automating task i is:

$$\text{Cost}_A(i) = \frac{r}{a(i)}, \tag{2}$$

with r denoting the rental rate of capital.[2] While automation decisions determine which tasks are automated, they may also affect the relative importance of different non-automated tasks

[1] Tasks are theoretical objects that structure production and are not directly observed in the data. The empirical analysis observes occupation-level skill requirements; the model therefore, derives occupation-level skill demand objects from this latent task structure.

[2] Following Acemoglu and Restrepo (2020), we assume perfect substitution between labour and capital at the task level, so each task is either fully automated or fully performed by labour. Across tasks, production aggregates in a Leontief manner, so that each task is essential and cannot be substituted by increasing input into other tasks.

within occupations. A task is automated if $\text{Cost}_A(i) < \text{Cost}_L(i)$.[3]

Let $\mathcal{A}$ denote the set of automated tasks and $\mathcal{L} = [0,1] \setminus \mathcal{A}$ the set of labour tasks. We define skill demand at the occupation level as the task-weighted importance of skill k within occupation o:

$$S_o^k \equiv \int_{i \in \mathcal{L}} s_k(i)\, \gamma_o(i)\, \mathrm{d}i. \tag{3}$$

This expression captures the share of tasks in occupation o performed by labour that require skill k, averaged over the set of non-automated tasks and weighted by the task distribution $\gamma_o(i)$. Automation decisions are task-specific and therefore common across occupations. Heterogeneity in the impact of automation arises from differences in task weights $\gamma_o(i)$ across occupations.[4]

2.2 Automation and Task Reallocation

Automation alters skill demand through two distinct channels. First, tasks in $\mathcal{A}$ are removed from labour input, which mechanically changes the composition of remaining tasks: skills concentrated in automated tasks lose weight in the occupation-level skill share simply because those tasks are no longer performed by workers. Second, the relative importance of surviving tasks may shift toward activities that are complementary to automated processes.

We capture this second channel by allowing automation to induce reallocation of effort across surviving tasks. Let $m(i) \in \{0,1\}$ indicate whether task i is complementary to automation, in the sense that its importance increases with the adoption or operation of automated processes (e.g., monitoring, coordination, or integration tasks). Occupations may differ in how strongly they reallocate effort toward such activities. We capture this possibility using an occupation-specific parameter $\theta_o \geq 0$, which summarizes the extent to which automation

[3]The analysis is in partial equilibrium: skill prices v_k and the rental rate of capital r are taken as given. In a full general equilibrium setting, both would adjust endogenously to the degree of automation and to the reallocation of labour and capital across sectors.

[4]Empirically, we measure skill demand at the occupation level rather than at the level of individual tasks. Our vacancy-based measure, the share of mentions of skill k over total skill mentions in postings for occupation o, provides a natural empirical counterpart to the occupation-level skill share S_o^k.

shifts effort toward complementary tasks within occupation o.

Post-automation task weights within occupation o are therefore defined as

$$\gamma'_o(i) = \frac{\gamma_o(i) \cdot \exp\big(\theta_o\, m(i)\big) \cdot \mathbf{1}\{i \in \mathcal{L}\}}{\int_{j \in \mathcal{L}} \gamma_o(j) \cdot \exp\big(\theta_o\, m(j)\big)\, \mathrm{d}j}, \qquad \theta_o \geq 0. \tag{4}$$

We adopt this exponential reweighting for analytical convenience, which ensures that task weights remain positive and sum to one. When $\theta_o = 0$, post-automation task weights correspond to pure truncation and renormalization of the task line. For $\theta_o > 0$, complementary tasks receive relatively greater weight among the surviving labour tasks.

We do not model the microfoundations of this reallocation explicitly. Instead, θ_o captures occupational features that may reflect organizational responses, production process characteristics, or complementarities between labour and automated capital (Autor, 2015; Acemoglu and Restrepo, 2022). This formulation allows the model to generate predictions about how automation affects within-occupation skill demand, conditional on the distribution of tasks and skills, without imposing a specific structure on the sources of task reweighting.

2.3 Decomposition of Aggregate Skill Demand

To study how automation affects skill demand, we introduce time variation and distinguish between within- and between-occupation adjustments. Let $t \in \{t_0, t_1\}$ denote the pre- and post-automation periods. Let $S^k_{o,t}$ denote the time-varying counterpart of the occupation-level skill share defined above in equation (3).

Aggregate skill demand at time t is defined at the level of the labour market as the employment-weighted average of occupation-level skill shares:

$$S^k_t = \sum_o w_{o,t} S^k_{o,t}, \tag{5}$$

where $w_{o,t} \equiv L_{o,t} / \sum_o L_{o,t}$ denotes the employment share of occupation o at time t.

Changes in aggregate skill demand can arise through two distinct margins: changes in the composition of employment across occupations, and changes in task composition within occupations. To make this distinction explicit, define the change in within-occupation skill shares as $\Delta S_o^k \equiv S_{o,t_1}^k - S_{o,t_0}^k$, the change in employment shares as $\Delta w_o \equiv w_{o,t_1} - w_{o,t_0}$, and the change in aggregate skill demand as $\Delta S^k \equiv S_{t_1}^k - S_{t_0}^k$. Expanding equation (5) across periods yields:[5]

$$\Delta S^k = \underbrace{\sum_o \Delta w_o \, S_{o,t_0}^k}_{\text{Between-Occupation Change}} + \underbrace{\sum_o w_{o,t_0} \, \Delta S_o^k}_{\text{Within-Occupation Change}} + \underbrace{\sum_o \Delta S_o^k \Delta w_o}_{\text{Interaction}}. \tag{6}$$

The between-occupation component captures changes in aggregate skill demand driven by shifts in employment shares across occupations, holding within-occupation skill shares fixed at their pre-automation values. The within-occupation component captures changes in task composition within occupations, holding employment shares fixed. The interaction term reflects the joint contribution of changes in both margins.

2.4 Implications and Testable Predictions

Equation (6) highlights that aggregate changes in skill demand arise through two distinct margins: occupational reallocation and within-occupation changes in task composition.

The former margin, i.e. shifts in occupational employment shares, is central to the canonical literature on skill-biased technical change and occupational polarization (e.g., Autor et al., 2003; Acemoglu and Autor, 2011; Goos et al., 2014), where changes in skill demand are primarily driven by the expansion or contraction of occupations.

In contrast, while task-based frameworks emphasize that technological change operates at the task level, the implications for within-occupation changes in multidimensional skill demand have received less direct emphasis in the literature. In our model, as formalized in

[5] See Section B in the Appendix for the derivation of equation (6).

equation (4), automation may alter the distribution of tasks performed within occupations, thereby reshaping occupation-level skill shares even when occupational employment weights remain stable.

For this reason, the theoretical results below focus primarily on the within-occupation margin. The model delivers predictions about how automation-induced task displacement and reweighting affect multidimensional skill demand within occupations, and how these changes translate into aggregate patterns.[6] Formal proofs of the propositions are provided in Appendix B.

(i) Skill Displacement Within Occupations

Proposition 1 (Automation-Driven Skill Displacement) *Fix occupation o and skill k. Abstracting from the reweighting channel in equation (4), if automation expands the set of automated tasks $\mathcal{A}$ so that the labour task set $\mathcal{L}$ shrinks, then the within-occupation skill share S^k_{o,t_1} decreases if skill k is used disproportionately in tasks that become automated relative to the surviving task set.*

When tasks are removed from the labour task set, skills used in those tasks become less prevalent within the occupation because equation (3) aggregates skill usage only over tasks performed by labour. If skill k is concentrated in tasks that are automated, its within-occupation share declines after automation. For example, if highly automatable activities such as routine data entry or repetitive assembly are automated, skills concentrated in those activities decline in within-occupation share. Accordingly, occupations more exposed to automation should exhibit reductions in the shares of skills primarily used in highly automatable tasks.

[6] All predictions operate through automation-induced changes in task weights $\gamma_o(i) \to \gamma'_o(i)$ in equation (4), while task-level skill requirements $s_k(i)$ remain fixed.

(ii) Within-Occupation Reallocation Toward Complementary Tasks

Proposition 2 (Task Reweighting toward Complementary Activities) *For occupation o, if $\theta_o > 0$ in equation (4), then complementary tasks ($m(i) = 1$) receive relatively greater weight in $\gamma'_o(i)$ than in $\gamma_o(i)$. If skill k has a higher average intensity in complementary tasks than in the average task of occupation o, then it is possible that $S^k_{o,t_1} > S^k_{o,t_0}$, even if occupation o does not perform tasks that are directly automated.*

A positive θ_o shifts task weights toward activities complementary to automated processes. Importantly, this mechanism operates even in the absence of direct automation of tasks within occupation o. If complementary tasks rely more heavily on skill k than the average task, reweighting increases the within-occupation skill share $S^k_{o,t}$. Empirically, automation exposure may therefore be associated with rising within-occupation shares of interpersonal, cognitive, or complex technical skills in occupations where complementary activities become more central.

(iii) Relative Skill Demand Within Occupations

While the previous propositions characterize changes in the level of within-occupation demand for individual skills, the following Proposition focuses on changes in the relative importance of different skills within occupations.

Proposition 3 (Relative Skill Upgrading) *Consider two skills k and ℓ, and assume $S^k_{o,t_0} > 0$ and $S^k_{o,t_1} > 0$. Suppose automation shifts task weight toward tasks in which skill ℓ is relatively more important than skill k, in the sense that tasks with lower automatability or complementary tasks ($m(i) = 1$) have weakly higher relative intensity in ℓ than in k. Then automation weakly increases the relative within-occupation demand for ℓ:*

$$\frac{S^\ell_{o,t_1}}{S^k_{o,t_1}} \geq \frac{S^\ell_{o,t_0}}{S^k_{o,t_0}}.$$

Automation reduces the weight of highly automatable tasks and may increase the weight of complementary tasks. Skills concentrated in less automatable or complementary activities therefore become relatively more prevalent within occupations. If skill ℓ is more strongly associated with such tasks than skill k, the relative within-occupation skill demand ratio $S^{\ell}_{o,t}/S^{k}_{o,t}$ increases after automation. Empirically, this implies that automation should raise the relative demand for skills associated with complementary or less automatable tasks compared to skills concentrated in highly automatable tasks. This mechanism is particularly relevant for occupations such as managers or professionals, where automation may not directly replace core tasks but may nevertheless raise the importance of coordination, oversight, and integration activities within the occupation.

(iv) Aggregate Skill Demand and Decomposition

The decomposition in equation (6) implies that aggregate changes in skill demand reflect the joint contribution of occupational reallocation and within-occupation changes in task composition.

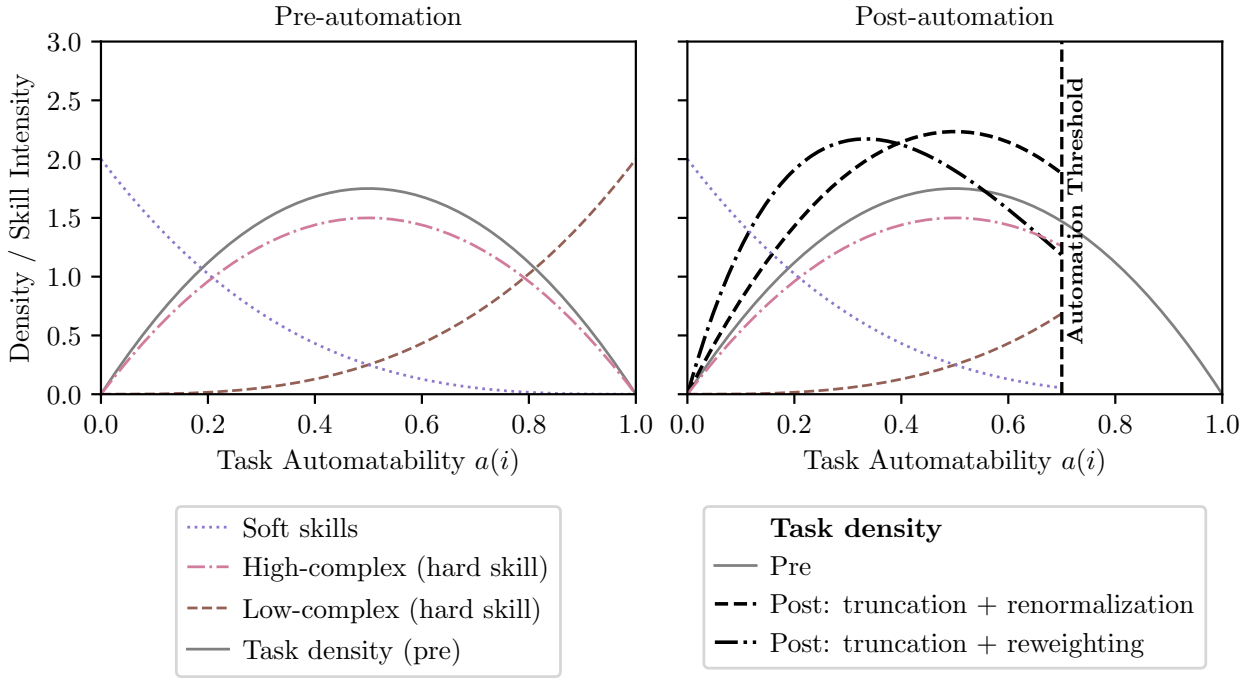
In particular, if changes in occupational employment shares Δw_o are small relative to within-occupation changes ΔS^{k}_{o}, then aggregate adjustment is primarily driven by within-occupation shifts in skill composition. Within-occupation changes are large when automation affects a substantial share of tasks or induces strong reweighting toward complementary activities, as captured by the change in task weights (see equation (4)). Whether aggregate skill demand changes are dominated by within- or between-occupation adjustments is therefore an empirical question, which we assess descriptively using the decomposition in Section 3.5, and the identification strategy in Section 4.

2.5 Illustrative Mechanisms

Figures 1 and 2 provide stylized illustrations of Propositions 1–3. Tasks are latent in the model and not observed empirically; the figures visualize how changes in task weights map

into changes in within-occupation skill shares. By construction, they abstract from changes in occupational employment weights and therefore do not depict the aggregate decomposition in equation (6), which is assessed empirically.

Figure 1: Task and Skill Distribution Before and After Automation

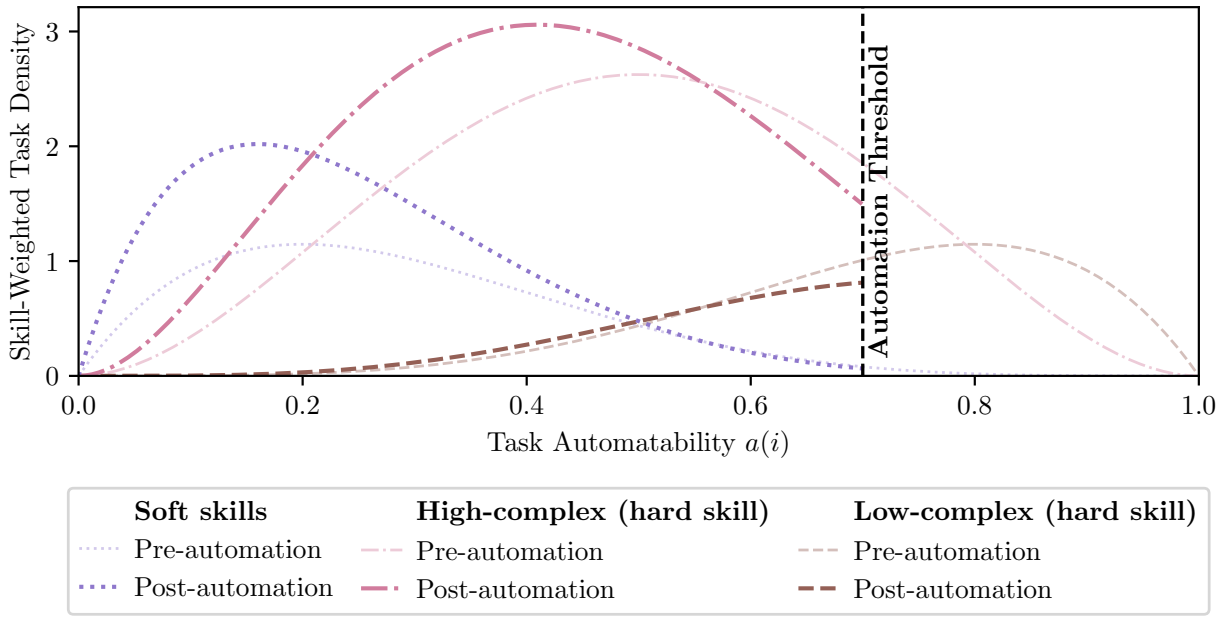


Notes: The left-hand panel illustrates a stylized pre-automation distribution of tasks within a representative occupation. The solid dark-gray curve (shown in both panels) captures the task density $\gamma_o(a)$ over the automatability index $a(i) \in [0, 1]$ and serves as a common baseline for comparison. Overlaid are three stylized skill-intensity profiles—low-complex hard skills, high-complex hard skills, and soft skills—representing the conditional likelihood $\mathbb{E}[s_k(i) \mid a(i) = a]$. These curves describe fixed task-level skill requirements. The right-hand panel shows post-automation task densities under two cases. Tasks with automatability $a(i) > 0.7$ are removed from labour, corresponding to automation adoption. Among the remaining tasks, the dashed black curve shows the post-automation task density under pure truncation and renormalization ($\theta_o = 0$), while the dash–dotted black curve allows for additional reweighting toward tasks complementary to automation ($\theta_o > 0$), as formalized in equation (4). Skill-intensity profiles are unchanged and are shown only over the surviving task domain. The vertical dashed line indicates the automation threshold.

Figure 1 illustrates how automation reshapes the distribution of tasks performed by labour through both displacement and reweighting. The left-hand panel shows the pre-automation task distribution together with skill-intensity profiles across the automatability dimension. The shape and relative positioning of these profiles are central for the model's predictions: skills that are concentrated in highly automatable tasks are more exposed to displacement, whereas skills concentrated in less automatable or complementary tasks become relatively more important as automation expands. Importantly, the theoretical predictions do not depend on the specific functional forms of these skill-intensity profiles, but only on the relative

concentration of skills in tasks with different levels of automatability or complementarity. The figures therefore provide a stylized illustration of the mechanisms rather than a direct mapping from observed data on task-level skill use.[7] The right-hand panel shows how these mechanisms operate. Pure truncation ($\theta_o = 0$) captures the displacement of highly automatable tasks (Proposition 1), while additional reweighting toward complementary tasks ($\theta_o > 0$) captures within-occupation shifts in task composition (Proposition 2).

Figure 2: Distribution of Skill-Weighted Task Density: Pre vs Post-Automation



Notes: The figure illustrates how changes in task weights translate into shifts in within-occupation skill demand. Each curve represents the task density weighted by skill intensity, $s_k(i)\gamma_o(i)$ (lighter colors) and $s_k(i)\gamma'_o(i)$ (darker colors), for low-complex hard skills, high-complex hard skills, and soft skills. Tasks are ordered along the horizontal axis by their automatability level $a(i) \in [0, 1]$. Post-automation task weights $\gamma'_o(i)$ reflect both (i) the removal of highly automatable tasks with $a(i) > 0.7$ and (ii) reweighting toward complementary tasks, as defined in equation (4). Task-level skill requirements $s_k(i)$ are held fixed. The area under each curve corresponds to the within-occupation skill share $S^k_{o,t}$. The figure abstracts from occupational employment changes and highlights within-occupation adjustments in task composition induced by automation.

While Figure 1 illustrates how automation alters the distribution of tasks within occupations, Figure 2 shows how these changes translate into shifts in within-occupation skill shares. Each curve plots the task density weighted by skill intensity, so that the area under the curve corresponds to the occupation-level skill share in equation (3). The figure thus illustrates how changes in task weights translate into the occupation-level skill shares that we estimate

[7] In the stylized example, highly automatable tasks are disproportionately associated with low-complex hard skills, making these skills more exposed to displacement. The same logic applies to any occupation in which certain skills are concentrated in tasks with higher automatability.

empirically. In particular, it illustrates how both task displacement (Proposition 1) and reweighting toward complementary tasks (Proposition 2) can generate relative increases in skills concentrated in less automatable or complementary tasks (Proposition 3).

3 Data

We use balance sheet data on machinery investment as our source of information on automation, combined with skill requirements extracted from job vacancy texts, covering the period 2008–2021. For identifying automation, balance sheet data offer an important advantage: they capture all final users of the equipment, whereas measures based on import flows miss machinery produced domestically or adopted through specialized integrator companies.

Balance sheet data are available for all sectors. However, some sectors record no machinery investment and are therefore excluded from the analysis. We further exclude the financial sector, as well as sectors in which machinery investment is more likely to reflect infrastructure-related capital (e.g., pumps or heating systems) rather than automation.[8] A detailed description of the sample restrictions is provided in Section C in the Appendix. We complement the main data with linked firm-level employment records from the FIRM register, occupation and worker background information from the IDAN and IDAP registers, import data from the UHDI register, and data on retraining courses from the VEUV register.

A central part of our analysis of within-occupation changes in skill demand focuses on heterogeneity in the effects of automation across occupations. To study this dimension, we distinguish between three broad occupational groups, referred to throughout as occupational groups. These groups correspond to three broad levels of work complexity: managers & pro-

[8]The final sample includes the following sectors: manufacturing, construction, retail, logistics, accommodation, publishing, real estate, professional activities, and arts. Automation investment is observed in all of these sectors, although its prevalence varies across them (see Table A1 in the Appendix). Table A2 in the Appendix reports the full list of included and excluded sectors.

fessionals, administrative & service workers, and production & elementary workers.[9]

This section describes how automation events are identified, characterizes the firms that adopt automation, and validates the approach. It then outlines the procedure used to extract skill requirements from job vacancy texts, presents descriptive evidence on skill demand in automating firms, and discusses the representativeness of the vacancy data. The section concludes by introducing the complementary register data used in the main analysis. Table A3 reports descriptive statistics for all variables introduced in the sections that follow. Finally, Section 3.5 presents a decomposition of skill demand among automating firms and discusses its implications.

3.1 Automation Events

Recent literature has exploited firms' tendency to adopt automation technologies through short periods of unusually large investment. This approach uses investment spikes to identify both which firms adopt automation and when adoption occurs (Bessen et al., 2023; Aghion et al., 2023; Hirvonen et al., 2025; Domini et al., 2021). We follow a similar strategy and identify lumpy machinery investments using balance sheet data from the FIRE register for Danish firms. These data cover physical assets acquired for use in firms' production of goods and services.[10] To identify significant spikes in yearly machinery investment, I_{ft}, we follow the approach used to identify robot adoption events in Békés et al. (2025). The automation event for firm f is defined as follows:

[9]The occupational groups are defined as follows: Managers & professionals, ISCO occupational codes 1, 2, and 3; Administrative & service, ISCO codes 4 and 5; Production & elementary, ISCO codes 6, 7, 8, and 9. Occupational heterogeneity in the effects of automation is examined by Acemoglu and Restrepo (2018) and is central to the literature on technology and work polarization (Autor and Dorn, 2013; Goos et al., 2014). Empirically, the polarization literature has adopted similar complexity-based occupational groupings to the one used here; see, for example, Domini et al. (2021).

[10]For this relatively broad definition to capture automation appropriately, the investment spikes must primarily reflect capital with autonomous functions, such as industrial robots and CNC machines. We verify that this is the case in Section 3.2.

$$\text{Aut}_{ft} = \begin{cases} 1 & \text{if } I_{ft} = \max_\tau I_{f\tau} \\ & \text{and } I_{ft} > 4 * \bar{I}_f \\ & \text{and } I_{ft} > 15,000 \text{ EUR } (112,000\text{DKK}) \\ 0 & \text{otherwise} \end{cases} \tag{7}$$

The criteria defined in Equation (7) imply that a firm-year is classified as containing an automation event if I_{ft} exceeds four times the firm's average level of machinery investment.[11] This ensures that the identified event represents an investment substantially larger than what is typical for the firm. Because some firms report no investment in most years, this criterion alone could classify a relatively small investment as an automation event. To avoid this, we additionally impose a minimum investment threshold of 15,000 EUR (112,000 DKK), below which no investment is classified as an automation event. We define event time relative to the automation event as $\tau \in \mathbb{Z}$, where $\tau = 0$ denotes the year of automation.

3.2 Automating Firms and Validation

As shown in Table 1, slightly less than one-third of firms (31%) experience an automation event. This figure is comparable to the 29% reported by Bessen et al. (2023), which is based on survey data on automation investment. Studies that identify automation spikes using import data typically report substantially lower shares of automating firms.[12] This difference is likely due to the fact that import data do not capture investments supplied by integrators or domestic producers. The share of automating firms in our data is also higher

[11] This threshold corresponds to a firm recording three equally sized investment episodes over the 14-year sample period from 2008 to 2021. To address the arbitrariness of the cutoff, we assess its robustness. As shown in Figure A1 in the Appendix, lowering the threshold to three increases the sample of automating firms by 23% (adding 16% of eligible firms), while raising it to five reduces the sample by 21% (excluding 15% of eligible firms).

[12] For example, Domini et al. (2021) identify automation events in 5–7% of firms.

than that typically reported in studies focusing specifically on industrial robots.[13] This is unsurprising given our broader definition of automation, which encompasses all machinery investment rather than investments in robots alone.

TABLE 1
Descriptive Statistics for Automating and Non-Automating Firms

	Automating Firms	Other Firms
No. Firms	4,982 (31%)	11,307
No. Firm-year obs	57,833	116,732
Avg. Revenue	86,967	65,641
Avg. Employment	53.3	60.8
Avg. Hourly wage	237.7	228.9
Avg. no. vacancies	19.9	27.0
Avg. no. skills per vacancy	3.1	2.9

Notes: This table reports descriptive statistics comparing firms that experience an automation event over the sample period 2008–2021 with firms that do not. An automation event is defined using the investment-spike criterion. Revenue, employment, and hourly wage are computed as the average over the yearly panel of firm-year observations. Revenue is reported in thousands of DKK; hourly wage is reported in DKK. Vacancies and skills per vacancy are based on the total posting of all firm-years in the sample (the yearly average of firm vacancies is 1.7 for automating firms). See Table A3 in the Appendix for full descriptive statistics on automating firms.

Table 1 also shows that there are notable differences between firms that experience automation events and those that do not. Firms with automation events have higher revenues than non-automating firms, but they are smaller in terms of employment. These differences make non-automating firms a less convincing control group. Accordingly, the remainder of the paper focuses exclusively on automating firms and identifies the effects of automation by exploiting variation in the timing of adoption, using firms that have not yet automated as the comparison group.[14]

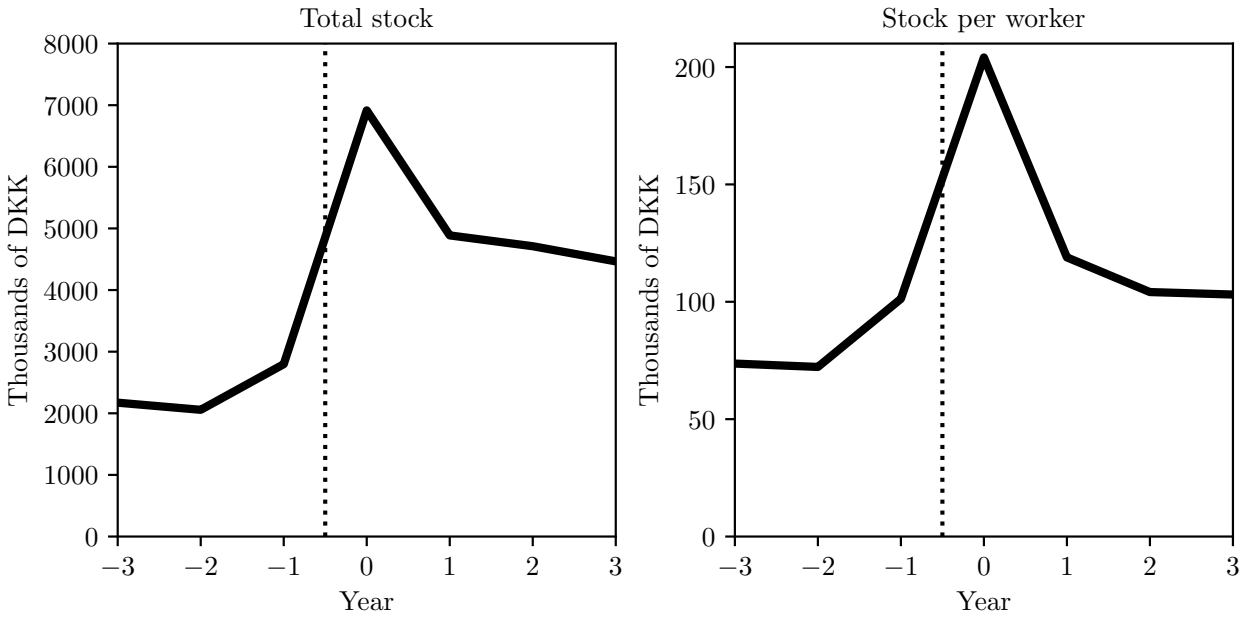
To illustrate the magnitude of the identified automation events, Figure 3 plots the evolution of machinery stock around the year of automation and shows a sharp increase at $\tau = 0$. On average, an automation event corresponds to a persistent doubling of machinery capital,

[13] For example, 1.5% of manufacturing firms in France (Bonfiglioli et al., 2024), 11.1% of manufacturing plants in the United States (Brynjolfsson et al., 2023), and between 1.1% and 12.3% of firms in Hungary, depending on the sample, (Békés et al., 2025).

[14] A similar identification strategy has been used in the automation literature, for example by Bessen et al. (2023).

with the stock increasing by 106% between $\tau = -3$ and $\tau = 3$, and to a 40% increase in machinery capital per worker.

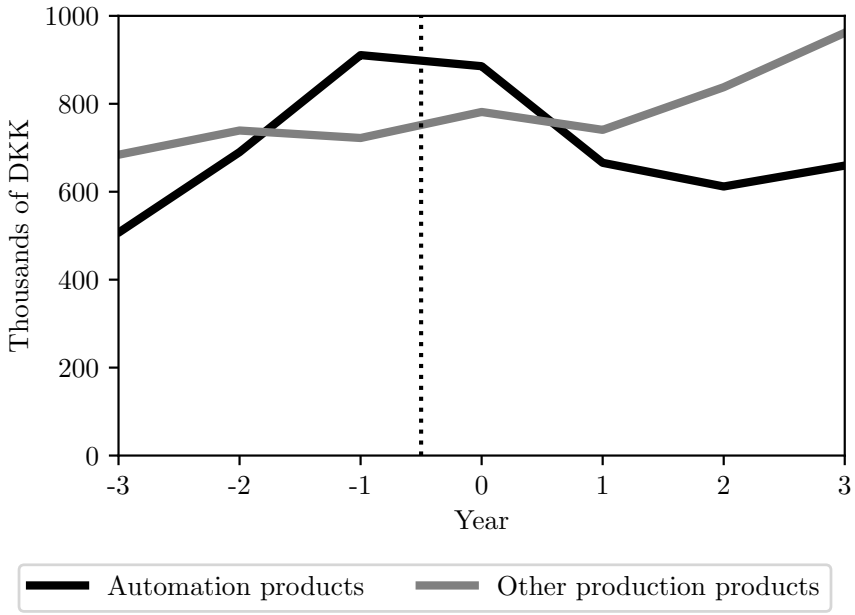
Figure 3: Machinery Capital Stock Around Automation Events



Notes: This figure displays the average machinery capital stock (left panel) and the average machinery capital stock per worker (right panel) of automating firms, plotted against event time $\tau \in \{-3, \ldots, 3\}$, where $\tau = 0$ denotes the year of the automation event as defined in Equation (7).

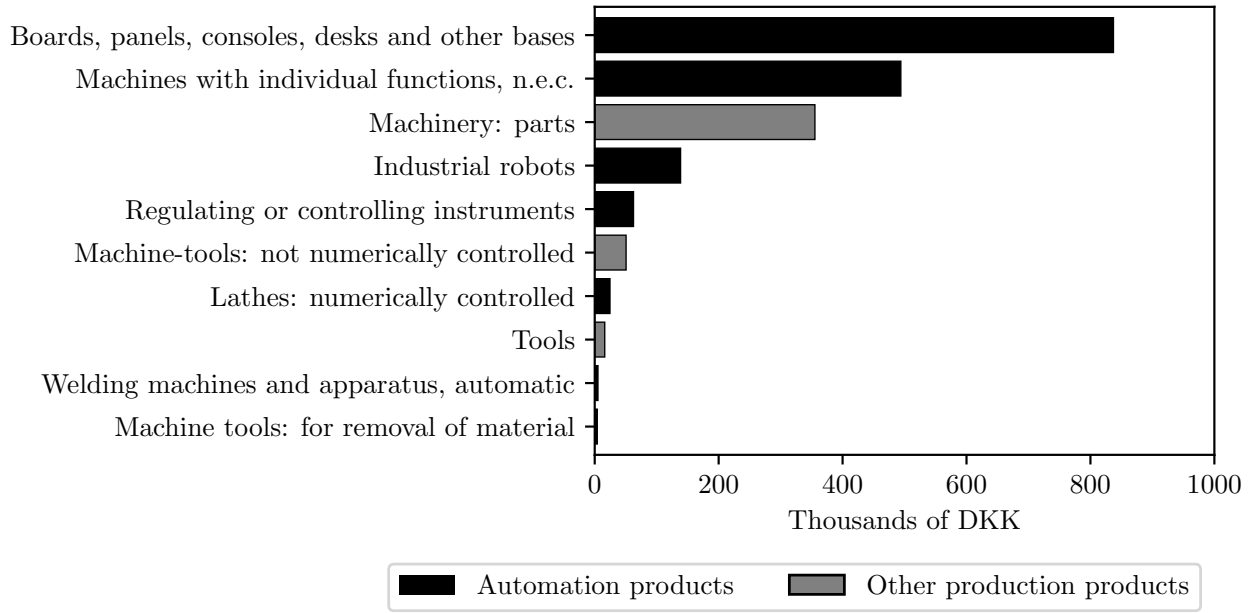
To validate that our broad measure of machinery investment captures automation technologies, we use firm-level import data on manufacturing equipment from the UHDI register to examine the type of machinery underlying these investments. Applying the automation classification of Acemoglu and Restrepo (2021), we distinguish between production goods with autonomous functions and other production goods. Figure 4 shows that the investment spike is driven primarily by products classified as having autonomous functions, while Figure 5 shows that imports around the identified automation year are dominated by such equipment. Taken together, these patterns provide reassuring evidence that the balance sheet measure captures investment in automation rather than capital investment more broadly.

Figure 4: Imports of Production Equipment Around Automation Events



Notes: This figure displays the average value of imports of production equipment around automation events, plotted against event time $\tau \in \{-3, \ldots, 3\}$, where $\tau = 0$ is the year of the automation event as defined in Equation (7). Import values are reported in thousands of DKK. Products are divided into automation products (products with autonomous functions) and other production products, following the classification of Acemoglu and Restrepo (2021).

Figure 5: Top 10 Production Equipment Imports



Notes: This figure displays the ten most imported categories of production products in the post-automation period $\tau \in \{1, 2, 3\}$, where $\tau = 0$ is the year of the automation event as defined in Equation (7). Import values are reported in thousands of DKK. Products are divided into automation products (products with autonomous functions) and other production products, following the classification of Acemoglu and Restrepo (2021).

3.3 Skill Demand

We use detailed job vacancy data to infer firms' skill demand. The vacancies are drawn from Jobindex, one of the largest online job platforms in Denmark.[15] The raw vacancy data were provided by a data consultancy company and consist of the universe of vacancies posted between 2008 and 2021. The final sample contains 58,728 vacancies posted by automating firms over the period 2010–2021.[16] From the vacancy texts, we extract information on both skills and occupations.[17]

We classify skill requirements into the three broad skill groups introduced in the conceptual framework in Section 2: soft skills, high-complex hard skills, and low-complex hard skills. These broad groups are constructed from the underlying skill categories developed by Deming and Noray (2020) and Deming and Kahn (2017).[18]

For the classification of vacancy skills, we define soft skills as general-purpose skills that are typically not acquired through formal education. Hard skills, by contrast, are more occupation-specific and may be acquired through formal education. We further distinguish between high- and low-complexity hard skills based on whether the skill is typically acquired through tertiary education. Figure 6 shows the three broad skill groups and their underlying skill categories, following Deming and Kahn (2017).

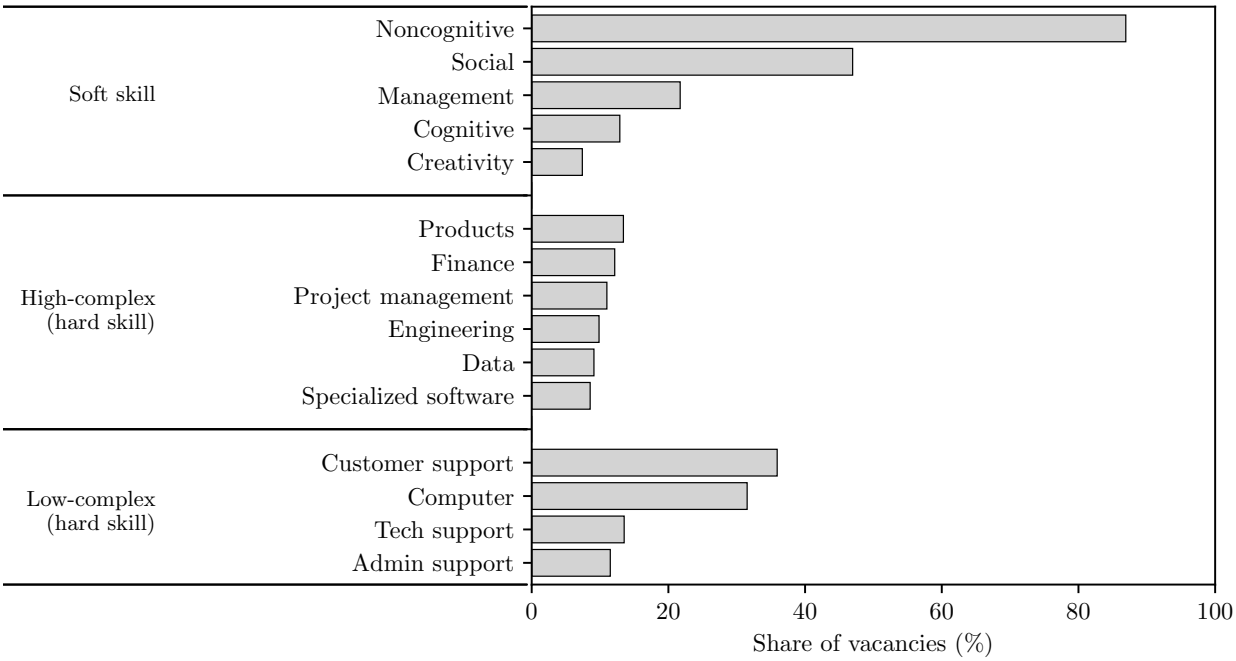
[15] Jobindex describes itself as the largest job platform in Denmark. Over the period 2008–2021, the data contain 0.16 posted vacancies for every individual hired in the Danish economy, compared with 0.65 in the final sample of automating firms. At the time of writing, Jobindex lists 34,400 active vacancies, compared with 21,600 on LinkedIn and 20,700 on the public employment service platform *Jobnet*.

[16] The full dataset contains 3.4 million vacancies. Most exclusions arise from restricting the sample to automating firms over the period 2010–2021, together with the additional sample restrictions described in Section C in the Appendix. This includes dropping most vacancies posted by public sector institutions, which are overrepresented in the raw data because of mandatory posting requirements, as well as vacancies without a formal job title.

[17] Details on the extraction of occupations from vacancy texts are provided in Section 1 of the Online Appendix.

[18] Our starting point is the set of skill categories in Deming and Noray (2020), which has also been used in, among others, Acemoglu et al. (2022) and Alekseeva et al. (2021). To provide a clearer picture of how automation affects specific dimensions of skill demand, we merge some categories from Deming and Noray (2020) while preserving their thematic content. Where appropriate, we align these merged categories with the broader classifications in Deming and Kahn (2017), combining categories that capture similar types of skills. This yields 15 relevant skill categories, listed in Table A4 in the Appendix.

Figure 6: Skill Content of Job Vacancies



Notes: This figure displays the skill content of job vacancies posted by automating firms in the year before the automation event ($\tau = -1$). The horizontal axis reports the share of vacancies requiring each of the 15 individual skill categories, grouped into three broad skill types. Soft skills are defined as general-purpose skills not typically acquired through formal education; high-complex hard skills are occupation-specific skills typically acquired through tertiary education; low-complex hard skills are occupation-specific skills not requiring tertiary education.

Our procedure for classifying skills from job vacancy texts begins by extracting key phrases that describe skill requirements. We do so using Named Entity Recognition based on a BERT model (Devlin et al., 2019), which yields on average 20 extracted phrases per vacancy. This initial step ensures that the classification is based only on text elements that are relevant to skill requirements. We then match the extracted key phrases to the keyword corpus associated with each underlying skill category in Deming and Noray (2020).[19] Because the keywords in Deming and Noray (2020) are in English, whereas the vacancy texts are in Danish, a simple keyword-matching approach is not feasible. To address this issue without relying exclusively on translations, which may not fully capture the expressions used by Danish recruiters, we use a sentence-embedding approach to measure semantic similarity between phrases.[20]

[19]For the full keyword corpus used in the classification, see Table 1 in the Online Appendix.

[20]This feature of the method is also advantageous in an English-to-English setting, as it allows phrases that are closely related to a skill category, but not exact lexical matches, to be classified appropriately.

We measure semantic similarity between phrases using pretrained sentence embeddings. Specifically, each phrase is encoded using the `distiluse-base-multilingual-cased-v2` model from the Sentence-Transformers framework.[21]

For each phrase s, the model produces a vector embedding $\mathbf{v}_s \in \mathbb{R}^{512}$. Semantic similarity between phrases i and j is then computed using cosine similarity,

$$\text{Sim}(i,j) = \frac{\mathbf{v}_i^\top \mathbf{v}_j}{\|\mathbf{v}_i\| \, \|\mathbf{v}_j\|}.$$

In practice, this procedure yields a similarity vector between each extracted vacancy phrase and each keyword in the Deming and Noray (2020) taxonomy. To classify vacancy phrases, we compute for each skill category the mean cosine similarity of the five most similar keywords to the phrase. A phrase is then assigned to the category with the highest similarity score, provided that this score exceeds a lower-bound threshold designed to exclude phrases that do not fit any of the 15 categories. The threshold is calibrated so that the share of phrases classified into skill categories closely matches that obtained through manual coding.[22]

The most frequently matched phrases in each skill category are manually reviewed. If a phrase is deemed appropriate, it is added to the corresponding skill corpus. Some phrases are erroneously matched because they are semantically close to a category without representing an actual skill requirement (for example, the word 'colleagues' being close to the social-skills corpus). These phrases are instead added to a separate exclusion corpus to prevent similar misclassification in later iterations. In subsequent rounds, a phrase is classified only

[21] Sentence-Transformers extend Transformer-based language models such as BERT to generate semantically meaningful sentence-level embeddings (Reimers and Gurevych, 2019). This addresses the language mismatch between the English skill taxonomy and the Danish vacancy texts, since the embedding model is multilingual and trained using multilingual knowledge distillation to align sentence representations across languages (Reimers and Gurevych, 2020).

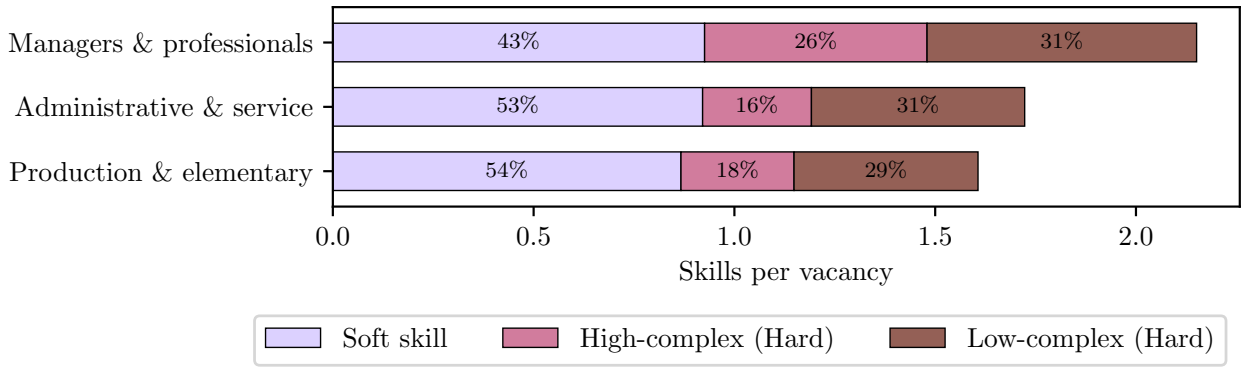
[22] Specifically, the 1,080 most common phrases were manually classified, of which 47% were assigned to a skill category; these manually classified phrases were also added to the skill-category corpus. We then iterated over increasing threshold values until the algorithm classified the same share, 47%, of phrases in this sample. The resulting threshold was subsequently applied to the remaining phrases. The manually classified sample was also used to evaluate the algorithm's accuracy, Accuracy $= \frac{TP+TN}{N}$, and to inform design choices such as the number of keywords against which each phrase is compared.

if its cosine similarity to the skill-category corpus exceeds its similarity to the corresponding exclusion corpus. After two rounds of manual review, the resulting corpus is used for the final classification of job vacancies. This procedure is not applied to software-related skill categories, since software names do not require translation or adaptation and are not handled well by the language model. Table A5 in the Appendix reports the most frequently classified keywords.

The resulting measure of vacancy skill requirements is a dummy variable indicating whether a vacancy requires a given skill category k. Consistent with the definition of occupational skill demand ($S^k_{o,t}$) in Equation 5 of the conceptual framework in Section 2, we express demand for skill k as its share in the total skills demanded within an occupation.

Figure 7 shows the composition of skill demand across the three broad occupational groups introduced at the beginning of Section 3. Managers & professionals unsurprisingly require more skills on average, and their relative skill composition also differs from that of less complex occupations. In particular, soft skills account for a smaller share of total skill demand in this group, while high-complex hard skills account for a larger share. By contrast, the difference in skill composition between administrative & service workers and production & elementary workers is comparatively small.

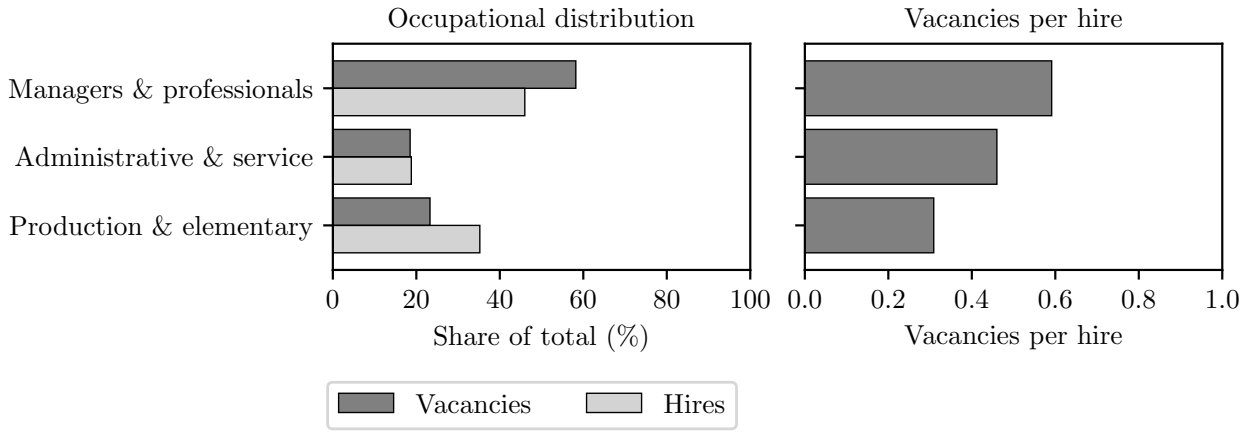
Figure 7: Skill Composition by Occupational Group



Notes: This figure displays the composition of skill demand across the three broad occupational groups, based on job vacancies posted by automating firms in the year before the automation event ($\tau = -1$). Occupational groups are defined by 1-digit ISCO-08 codes: managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). The horizontal axis reports the mean number of skills identified per vacancy. The percentage values within each bar report the share of total skill mentions belonging to each broad skill type (soft skills, high-complex hard skills, low-complex hard skills). These shares sum to 100% within each occupational group.

Finally, we assess the representativeness and relevance of the vacancy data and the skill measures derived from it. Figure 8 shows that vacancy postings among automating firms are skewed toward managers & professionals, who account for 60% of vacancies but 50% of hires (left-hand panel). This overrepresentation may reflect the greater use of online job platforms in these occupations, as well as some measurement error in assigning occupations to vacancy postings. By contrast, production & elementary workers are underrepresented, with roughly half as many posted vacancies per hire as managers & professionals (right-hand panel). This bias toward higher-skill occupations is common in vacancy data, with similar patterns documented, for example, by Atalay et al. (2020) and Hershbein and Kahn (2018).

Figure 8: Vacancies and Hires by Occupational Group



Notes: This figure compares hiring and vacancy posting patterns across the three broad occupational groups among automating firms over the period 2008–2021. Occupational groups are defined by 1-digit ISCO-08 codes: managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). The left-hand panel shows the share of total hires and the share of total vacancies accounted for by each occupational group. The right-hand panel shows the ratio of posted vacancies to hires for each group. For representativeness of more granular occupations, see Figure 1 in the Online Appendix.

In the context of this paper, the representativeness of vacancy counts is not a major concern, since our main focus is on within-occupation skill shares. Nevertheless, the composition of the vacancy data may still matter in two ways. First, managers & professionals are overrepresented in the vacancy data underlying the firm-level analysis. Second, variation is more limited for the somewhat underrepresented group of production & elementary workers, which may make it more difficult to detect effects for this group.

The relevance of the skill categories is another potential concern. Because the skill taxonomies developed by Deming and Noray (2020) and Deming and Kahn (2017) are primarily designed for professional occupations, it is not obvious that they are equally well suited to occupations of lower complexity. Figure 7 is nevertheless reassuring, as production & elementary occupations display only about 25% fewer skills per vacancy than managers & professionals. A related question is whether the skill measure captures skills that are plausibly affected by automation. Table A6 in the Appendix reports a sample of phrases with high similarity to production-related language. Among soft skills, these phrases relate mainly to the management and design of production. High-complex hard skills contain the richest set of production-related terms, ranging from specific production methods to craftsmanship. Low-complex hard skills contain fewer such terms, mostly related to machinery maintenance. This suggests that the measure does capture skills relevant to automation, although it also raises the possibility that low-complex skills related to operating and supervising machinery are not fully captured. Under the assumption that this measurement error is uncorrelated with automation events, any resulting omission would mainly reduce the precision of the estimates, particularly for production & elementary occupations.

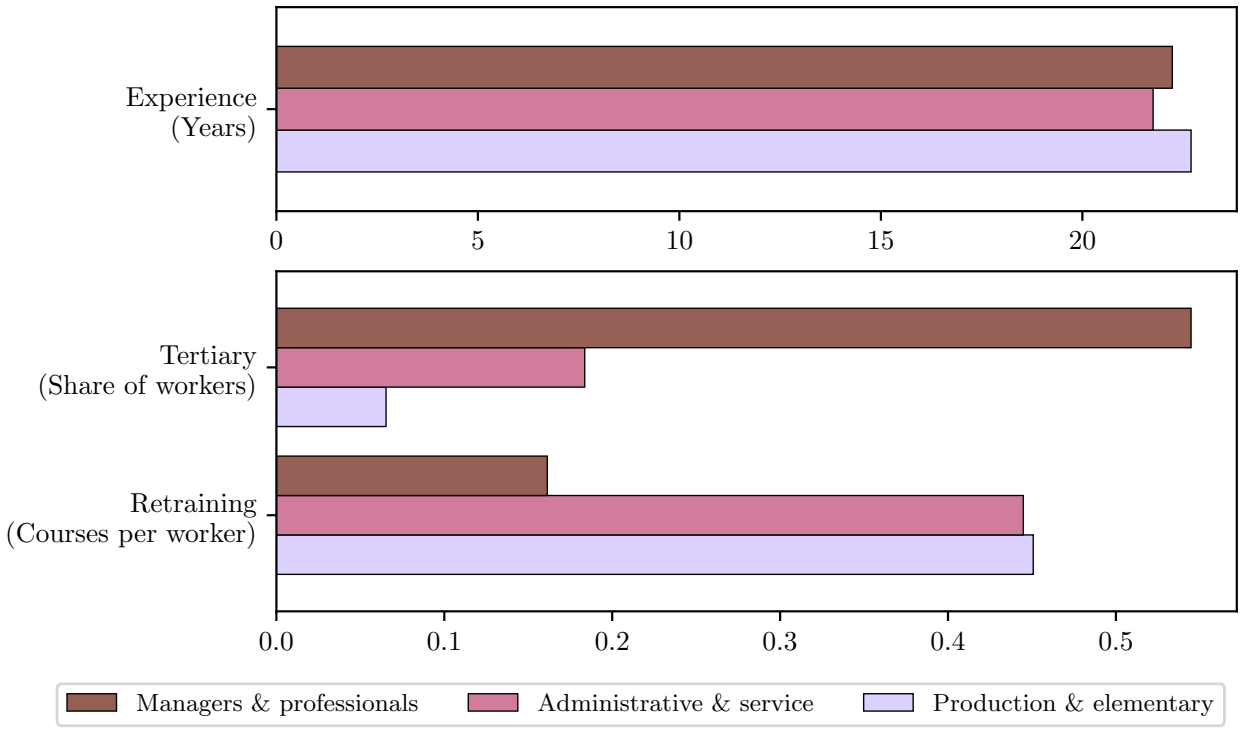
To address concerns about potential measurement error in the skill measures and the lower variation observed in some occupations, our analysis leverages complementary register-based proxies for worker skills, including retraining programs relevant to production & elementary workers. The next section discusses these measures.

3.4 Register data and complementary measures of skill demand

Job vacancy data provide a valuable source of revealed information on the skills demanded by firms. At the same time, they have some limitations, including occupational bias, the inability to show whether firms actually hired workers matching the posted requirements, the possible omission of skills taken for granted in a given position, and limited variation in a firm-level event-study setting.

To address these limitations, the empirical analysis complements the vacancy-based measures with register-based proxies for worker skills, namely worker experience, education, and participation in retraining courses.[23] Figure 9 shows the prevalence of these measures across occupational groups.

Figure 9: Complementary Skill Measures



Notes: This figure displays three register-based proxies for worker skills, measured among employees of automating firms in the year before the automation event ($\tau = -1$), separately by occupational group (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). *Experience* is the mean number of years of labour market experience of employees. *Tertiary* is the share of employees with any tertiary-level degree. *Retraining* is the number of adult education courses taken in the year $\tau = -1$, divided by the number of employees.

Worker experience is our closest register-based proxy for soft skills. We measure experience as the average number of years that a firm's workers have spent in the labour market prior to the current year, based on pension contribution records. This measure is most closely related to soft skills, which are more general-purpose and more likely to be accumulated through labour market experience than through formal education. As with soft skills in the vacancy data, average experience is relatively similar across occupational groups in absolute

[23]The Danish register data cover the entire working-age population. Information on experience and education comes from the IDAN and IDAP registers, while data on retraining are drawn from the VEUV register.

terms.

To proxy high-complex hard skills, we use the share of employees with a tertiary degree. This corresponds closely to our definition of high-complex hard skills as skills typically acquired through tertiary education. As expected, tertiary education is concentrated among managers & professionals. This mirrors the distribution of high-complex hard skills in the vacancy data, although the concentration in higher-skill occupations is even more pronounced for formal education.

Finally, we use data on participation in public retraining courses to capture the adjustment mechanism potentially involving low-complex hard skills. Denmark has an extensive adult education system, including courses aimed at reskilling both low- and high-skill workers. These courses, known as AMU (*arbejdsmarkedsuddannelser*, or labour market training), can be taken by both unemployed and employed workers.

TABLE 2
Most Common Retraining Courses by Occupational Group

Managers & professionals		Administrative & service		Production & elementary	
Course name	%	Course name	%	Course name	%
Commerce & administration (AMU)	12.4	Commerce & administration (AMU)	20.8	Transportation (AMU)	23.1
Management (AMU)	7.7	Metal industry (AMU)	18.9	Construction (AMU)	11.4
Metal industry (AMU)	5.9	Transportation (AMU)	7.8	Metal industry (AMU)	8.9
Management (Other)	5.7	Management (AMU)	5.0	Iron and metal industry (AMU)	8.3
Public Business Economics (Other)	4.6	Management (Other)	4.3	Welding (AMU)	5.0

Notes: This table reports the five most common publicly funded adult education courses taken by employees of automating firms in the year following the automation event ($\tau = 1$), separately by occupational group (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). Course data are drawn from administrative data that records all individual participation in publicly funded adult education courses in Denmark. The % column indicates the number of participants in a course as a share of total course participation within the relevant occupational group and period. Course names are translated to English by the authors. "AMU" denotes courses in the Danish labour market training system (*arbejdsmarkedsuddannelser*); "Other" denotes non-AMU adult education courses.

Courses are typically designed around specific work tasks, such as welding, food safety, and truck driving. In the context of this paper, they are particularly useful for capturing adjustment mechanisms that may be relevant for production & elementary workers, although some courses are also relevant for other skill dimensions and occupational groups, such as management. Table 2 reports the most common courses taken by employees in each occupational group.

The data on adult retraining come from the VEUV register, which records all individual participation in publicly funded adult education courses. The register contains information on course dates and subjects. Although we do not observe whether participation is initiated by the employer or by the worker, these data still provide useful complementary evidence on changing skill needs within firms. We define retraining as the number of courses taken relative to the number of workers in the firm. Retraining is observed across all occupational groups, but it is more common among lower-complexity occupations. The intensity of retraining is also substantial: in a given year, there is nearly one course for every two employees among production & elementary workers and among administrative & service workers.

Although the focus of the paper is on skill demand, a natural related question is how firms acquire and shed skills. We therefore also examine employment dynamics, including hiring and separations. For consistency, these measures are constructed using the number of individuals employed in the firm, where hires are defined as workers not employed by the firm in the previous year and separations as workers no longer employed by the firm in the following year.

3.5 Decomposition of Skill Demand

Propositions 1–3 in the conceptual framework predict how automation may generate within-occupation changes in skill demand. As discussed in point (iv) of Section 2.4, whether overall changes in skill demand are driven primarily by within-occupation adjustment or by changes in occupational employment shares remains an empirical question. In this section, we assess the relative importance of these channels through a descriptive decomposition of changes in skill demand. Section 4 then presents causal evidence on the effects of automation on the composition of skill demand.

We fit our vacancy data to the skill decomposition of the conceptual framework in Equation

(6) for each skill type k.[24] We construct pre- and post-automation periods by pooling vacancy postings in automating firms over the intervals $\tau \in \{-3, -2, -1\}$ (pre, denoted t_0) and $\tau \in \{1, 2, 3\}$ (post, denoted t_1).

Let S^k_{o,t_0} denote the share of skill k in the total number of skills demanded within the 2-digit ISCO-08 occupation o in the pre-automation period,[25] while w_{o,t_0} denotes the share of vacancies posted for occupation o in the same period. Analogous definitions apply for t_1, while Δ denotes the change between post- and pre-automation periods. Because our focus is on changes in skill composition rather than entry and exit from vacancy posting, we restrict the sample to firms with at least one job vacancy posted in both periods.

The Between-Occupation component in Equation (6) can be interpreted as the change in skill demand due to a change in the composition of occupations demanded by employers, while holding skill requirement within each occupation fixed at the baseline period. This component captures the importance of the adjustment in the demand for skills due to the change in demand for occupations with differential baseline skill requirements, assuming that the skill composition required to perform them remains as at baseline. The Within-Occupation term represents the variation of skill demand that is due to the change in skill shares within each occupation, keeping the occupational composition fixed at the baseline period. This component is meant to capture the extent to which skill requirements change over time for the same occupation. The Interaction component measures the part of the variation in total skill demand due to the interaction between changes in the occupational composition and changes in the skills demanded for a given occupation. The interaction term is positive when changes in occupational weight and changes in the share of skill demand within that occupation move in the same direction.[26]

[24] Similar decomposition approaches have been used in the prior literature in other contexts, e.g., Battisti et al. (2023), Petrova et al. (2024).

[25] Note that this definition of skill shares is invariant to aggregate trends in vacancy posting and in the frequency of skill mentions in vacancies.

[26] In other words, the term is positive when changes in between-occupation demand and changes in within-occupation skill demand are aligned. Conversely, the term is negative for a given skill k when occupations that increase their demand for skill k become less prominent among automating firms, or vice versa.

The aggregate decomposition of skill change may conceal substantial heterogeneity across occupational groups. A given component may appear small in magnitude because it is positive for one occupational group and negative for another, thereby offsetting in the aggregate. More generally, each component of skill change may be driven by a subset of occupations rather than being evenly distributed across them. We therefore further decompose each term on the right-hand-side of Equation (6) into changes for the three broad occupational groups: managers & professionals, administrative & service, and production & elementary occupations. For instance, the within component of change in demand for skill k can be further decomposed as follows:

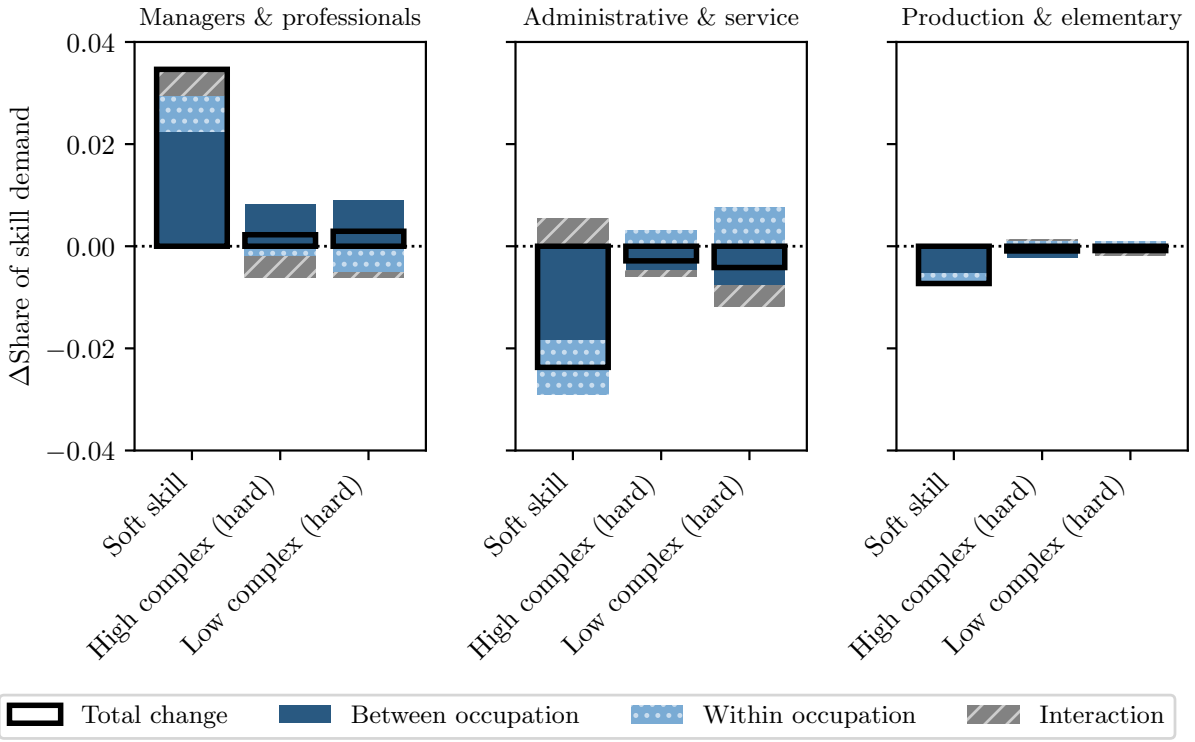
$$\Delta^{Within} S^k \equiv \sum_o \Delta S_o^k w_{o,t_0} = \sum_g \sum_{o \in g} \Delta S_o^k w_{o,t_0} \tag{8}$$

with g being one of the three occupational groups. The other two terms in equation (6) can be decomposed in a similar way.

Figure 10 presents the result of the descriptive decomposition, revealing heterogeneity across the occupational groups. Demand for soft skills increases for managers & professionals, driven by both an increase in vacancies in underlying 2-digit occupations that were more soft-skill intensive at baseline, as well as a rise in soft-skill intensity within those occupations. In contrast, for administrative & service occupations the alignment of the two components occurs in the opposite direction, showing a decrease in demand for occupations with relatively high baseline soft-skill intensity, and a reduction of their intensity within those vacancies. For production & elementary roles, the changes are minor and point to a decline in demand for soft skills, mainly driven by occupational reallocation.

The change in hard-skill demand displays a misalignment between the within- and between-occupation components. Among managers & professionals, automation is associated with increased demand for hard-skill-intensive occupations. In contrast, among administrative & service occupations, automation induces a reallocation away from hard-skill-intensive roles,

Figure 10: Decomposition of Skill Demand by Occupational Group



Notes: This figure displays the components of the descriptive decomposition of skill demand for automating firms, following Equation (6). Each bar reports the total change in skill demand along with its three components: the between-occupation component ($\sum_{o \in g} \Delta w_o \, S^k_{o,t_0}$), the within-occupation component ($\sum_{o \in g} w_{o,t_0} \, \Delta S^k_o$), and the interaction component ($\sum_{o \in g} \Delta S^k_o \, \Delta w_o$). The decomposition is computed using Equation (8)), which further disaggregates each component of Equation (6) by occupational group g. Pre- and post-automation periods are defined by pooling vacancies over $\tau \in \{-3, -2, -1\}$ and $\tau \in \{1, 2, 3\}$, respectively. Skill shares $S^k_{o,t}$ are measured at the two-digit ISCO-08 occupation level. Employment shares $w_{o,t}$ are measured as the share of vacancies posted for occupation o. Each panel corresponds to one of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). See Figure A2 in the Appendix for the overall firm-level decomposition, and Table 3 for the quantification of the within-occupation share of total adjustment.

while hard-skill intensity rises within the remaining occupations.[27]

These heterogeneous patterns across occupational groups, which cannot be detected using administrative data alone, are masked in the firm-level decomposition in Equation (6). Figure A2 in the Appendix reports the overall decomposition of skill change for automating firms. It shows that the relationship between automation and skill demand is not driven by uniform within-occupation or between-occupation shifts, but instead by sizable interaction components, consistent with joint movements in occupational restructuring and within-occupation changes in skill intensity that differ across occupational groups. This motivates

[27]This suggests that the types of hard skills required by managers & professionals differ from those required in administrative and clerical occupations, and may therefore be affected differently by automation.

the separate analysis by occupational group in the main empirical analysis of Section 4.

Overall, the decomposition exercise shows that automation is not simply skill-biased in the classical sense, but instead reshapes the demand for multidimensional skills in a more nuanced way. Automation reallocates demand toward managers & professionals who themselves are becoming more soft-skill intensive. At the same time, it reduces demand for routine administrative work, while the remaining administrative & service positions require fewer soft skills and more hard skills. Demand for production & elementary occupations also declines, although skill requirements within these roles remain broadly stable. This points to an increased demand for coordination, leadership, and complex decision-making skills within managerial occupations, alongside a reduction in the relative importance of administrative and production workers.

TABLE 3
Within-Occupation Contribution to Skill Adjustment

	Soft Skills	High-Complex (Hard)	Low-Complex (Hard)
Managers & professionals	20%	14%	34%
Administrative & service	31%	34%	39%
Production & elementary	27%	29%	35%

Notes: This table reports, for each occupational group and skill type, the share of total absolute skill demand adjustment accounted for by the within-occupation component of the decomposition in Equation (6). Specifically, for occupation group g the share is computed as follows: $\frac{|Within_g^k|}{|Between_g^k|+|Within_g^k|+|Interaction_g^k|}$, or $\frac{\left|\sum_{o\in g}\Delta S_o^k w_{o,t_0}\right|}{\left|\sum_{o\in g}\Delta w_o S_{o,t_0}^k\right|+\left|\sum_{o\in g}\Delta S_o^k w_{o,t_0}\right|+\left|\sum_{o\in g}\Delta S_o^k \Delta w_o\right|}$, for each skill type k. Occupational groups are defined by 1-digit ISCO-08 codes: managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9).

To provide some indication of the importance of the within-occupation component of skill changes, Table 3 reports the corresponding share in the total absolute skill demand adjustments. Within-occupation skill change is a sizable contributor to overall change, ranging from 14% for high-complex hard skills among managers & professionals to 39% for low-complex hard skills in administrative & service occupations.[28]

[28] As expected, these values are lower than the 85% within-occupation task change found by Spitz-Oener (2006) and the 66% within-occupation skill supply change found by Salomons et al. (2025). Both studies capture long-run trends over decades of technological shift, whereas our estimates reflect adjustment to discrete automation investments over a much shorter horizon.

The next section uses a difference-in-differences strategy to identify the causal effect of automation on skill demand within occupational groups and within detailed occupations.

4 Empirical Strategy & Results

In this section, we estimate the causal effects of automation on skill demand. We interpret the results through the lens of the conceptual framework, focusing on skill displacement (Proposition 1), reallocation toward complementary tasks (Proposition 2), and relative skill upgrading (Proposition 3), while also distinguishing between within- and between-occupation margins of adjustment.

We exploit the staggered adoption of automation across firms, as identified in Section 3.1, to recover a suitable comparison group for adopting firms. Since automating firms may differ from non-adopters along several dimensions, including size, productivity, and pre-existing skill mix (Bessen et al., 2023), a standard two-way fixed effects estimator would conflate the effect of automation with underlying heterogeneity across firms, and may produce biased estimates when treatment effects vary over time and already-treated firms are used as implicit controls (Goodman-Bacon, 2021). To address these challenges, we use the approach proposed in Callaway and Sant'Anna (2021) and restrict comparisons to firms that have not yet adopted automation by period $\tau = 0$. This allows us to avoid the "forbidden comparison" issue and bypasses the concern that adopting and non-adopting firms are structurally different, since we use the same firms in periods *prior* to their automation event as the comparison group, exploiting the within-firm variation in a staggered event-study design.

Given the substantial heterogeneity in the evolution of skill demand across occupational groups documented in Section 3.5, we conduct our primary analysis separately for each group g: managers & professionals, administrative & service occupations, and production &

elementary occupations.[29] We are interested in the Average Effect of Automation (AEA), defined as follows:

$$AEA^g(c,t) = \mathbb{E}\left[Y_{ft}^g(1) - Y_{ft}^g(0) \mid C_f = c\right].$$

where f indexes firms, g is each of the three occupational groups and t indexes year. C_f denotes the automation event year of firm f. Potential outcomes $Y_{ft}^g(1)$ and $Y_{ft}^g(0)$ denote the outcomes with and without automation events. Under the no-anticipation and parallel trends assumptions, the AEA can be expressed as a two-by-two difference-in-differences using not-yet-adopters as controls:

$$AEA^g(c,t) = \mathbb{E}\left[\left(Y_{ft}^g - Y_{f,c-1}^g\right) \mid C_f = c\right] - \mathbb{E}\left[\left(Y_{ft}^g - Y_{f,c-1}^g\right) \mid C_f > t\right], \qquad t \geq c,$$

with baseline period $c-1$. Y_{ft}^g denotes the firm-level outcome measuring the share of each skill type demanded by firm f in vacancies for occupational group g at time t. The first term is the pre-post difference for firms first adopting in year c, while the second term uses as comparison units those firms that have not yet adopted automation by year t. We estimate the average effects of automation using the two-step method proposed by Callaway and Sant'Anna (2021) and aggregate the individual cohort-time effects $\widehat{AEA^g}(c,t)$ into a single overall treatment effect:

$$\widehat{AEA^g_{post}} = \sum_c \sum_{t \geq c} w_{c,t}\, \widehat{AEA^g}(c,t), \tag{9}$$

where $w_{c,t}$ are weights proportional to the number of treated observations in cohort c at time t, normalized to sum to one across all post-treatment cohort-time pairs. This aggregation summarizes the average effect of automation across all post-event periods and, given the limited precision of period-by-period effects for some dependent variables, it is

[29] We also estimate the effect of automation on skill demand at the firm level. Consistent with the small aggregate effects documented in the descriptive decomposition in Section 3.5, the firm-level estimates are generally small and statistically insignificant, suggesting that the aggregate effect of automation on the skill mix masks substantial heterogeneity across occupational groups.

our preferred representation.[30] We report estimates $\widehat{AEA^g_{post}}$ with 95% bootstrap confidence intervals.

In addition to skill demand measures, we apply the same empirical strategy to two sets of complementary outcomes that help shed light on the mechanisms underlying workforce adjustments induced by automation. First, we consider administrative-data measures of firms' realized human capital, rather than the desired skill mix inferred from job postings. These include average labour market experience, the share of tertiary-educated workers, and retraining intensity. Second, we examine the evolution of employment and worker flows, including hiring and separations.

The parallel trends assumption requires that adopters and control units would have followed similar trajectories in the absence of automation. By using not-yet-adopters as controls and comparing firms that adopt automation at different points in time, we mitigate concerns about systematic differences between adopters and never-adopters. Identification requires that future adopters provide a valid counterfactual prior to their own automation. In addition, the no-anticipation assumption requires that automation does not affect outcomes before it occurs. Although these assumptions cannot be tested directly, we provide supporting evidence by inspecting event-study dynamics and show that there are no statistically significant pre-treatment effects or systematic divergence between adopters and not-yet-adopters prior to the automation event.[31]

It is worth noting that the group-level estimates $\widehat{AEA^g_{post}}$ capture the effect of automation on skill demand within each broad occupational group, but do not distinguish whether this effect operates through changes in the skill content of individual occupations or through

[30] In robustness Section 4.3, we also report dynamic estimates relative to the treatment year, where $\widehat{AEA^g}(c,t)$ are aggregated across adopting cohorts: $\widehat{AEA^g_s} = \sum_c w_c \, \widehat{AEA^g}(c, c+s)$,

[31] Our baseline estimates rely on an unconditional parallel trends assumption. We nonetheless report robustness checks that condition on a set of pre-treatment firm characteristics, including one-digit industry, firm-size quartiles and quartiles of the provincial share of tertiary-educated workers, the latter capturing potential differences in local skill supply. Reassuringly, the event-study estimates show no statistically significant pre-event dynamics both with and without controls.

reallocation across occupations within the group. To the extent that automation reshuffles the composition of vacancies across two-digit occupations belonging to the same group g, the group-level estimates will reflect both channels. To isolate the within-occupation component of the effect, we estimate an analogous event study at the level of individual two-digit occupations. Specifically, for each occupation o, we run the Callaway and Sant'Anna (2021) estimator on the subset of firms posting vacancies in o, where the outcome Y_{ft}^{o} measures the skill share demanded by firm f in occupation o at time t, yielding occupation-specific dynamic effects $\widehat{AEA}_{post}^{o}$.[32] This occupation-level analysis serves to identify which specific occupations experience genuine within-occupation changes in skill demand following automation.

4.1 Automation & Skill Demand

In this section we discuss results on the Average Effect of Automation on skill demand and workforce composition. Figure 11 presents the aggregate post-treatment estimates $\widehat{AEA}_{post}^{g}$ from equation (9), separately for the three occupational groups.

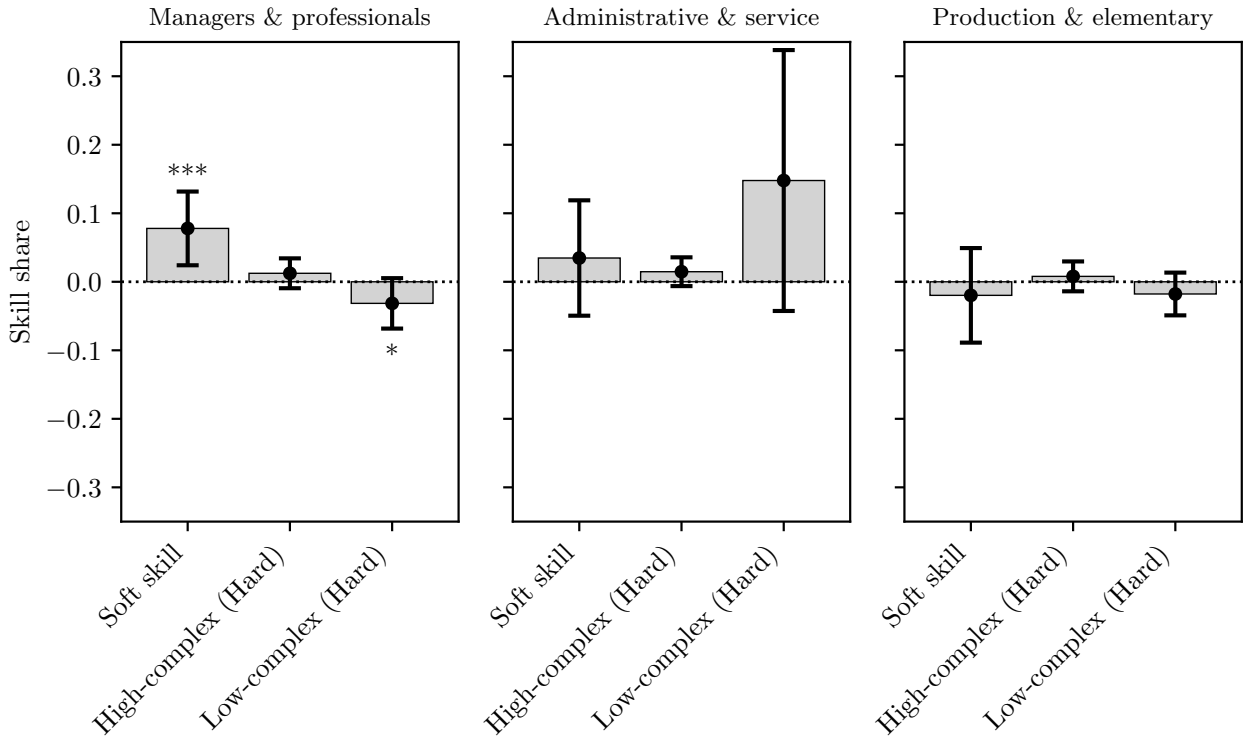
In line with the descriptive evidence of Figure 10, automation leads to a significant increase in the share of soft skills demanded in manager & professional occupations (approximately 8 percentage points, significant at the 1% level). Consistent with the within-occupation decrease in low-complex hard skills, the results also suggest a decline in the demand for low-complex hard skills (around 3 percentage points, significant at the 10% level), implying that automation crowds out low-complex technical requirements.[33]

The increase in soft skills in managers & professionals is consistent with the automation-task complementarity mechanism in Proposition 2. Even if few tasks in these occupations

[32] For each event-time s, the estimation sample includes only firms with at least one vacancy in occupation o in both the baseline period $c-1$ and in period $c+s$, so that each estimate captures changes in the skill content of a given occupation within continuing firm–occupation spells. Identification requires that the parallel trends and no-anticipation assumptions hold separately for each occupation: among firms posting in occupation o, adopters and not-yet-adopters must follow similar skill-demand trajectories absent automation.

[33] However, the result for low-complex hard skills among managers & professionals is sensitive to alternative definitions of automation events; see Section 4.3 for details.

Figure 11: Effects of Automation on Skill Demand by Occupational Group



Notes: This figure reports the aggregate post-treatment Average Effect of Automation ($\widehat{AEA}^{g}_{post}$) on skill demand, estimated using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021) as described in Equation (9). The estimator uses not-yet-adopter firms as the comparison group and aggregates cohort-time effects across all post-event periods, with weights proportional to the number of treated observations in each cohort-time cell. The outcome variable in each case is the share of a given broad skill type (soft skills, high-complex hard skills, or low-complex hard skills) in total skills demanded within the occupational group. All outcomes are weighted by the number of vacancies. Each panel displays estimates from a separate event study for one of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). The sample consists of 3,571 automating firms. The skill content is based on 33,747 vacancies (9,293 firm-year observations) for managers & professionals, 10,497 vacancies (3,480 obs.) for administrative & service, and 13,707 vacancies (5,899 obs.) for production & elementary. Full coefficient table reported in Table 2 of the Online Appendix. See Figure A3 for firm-level estimates and Figure A4 for dynamic event-study estimates. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

are directly automated, automation raises the importance of activities such as coordination, communication, and judgment, which rely more intensively on interpersonal and cognitive skills.[34] This interpretation aligns with evidence that technological change shifts work toward non-routine, interactive, and problem-solving activities (e.g., Autor, 2015; Deming, 2017; Acemoglu and Restrepo, 2022), increasing the relative importance of skills that complement new technologies. Furthermore, the co-occurrence of rising soft skills and declin-

[34]Through the lens of the conceptual framework in Section 2, $m(i) = 1$ for several tasks and $\theta_o > 0$ for managerial and professional occupations, implying that these occupations contain tasks that are strongly complementary to automated processes.

ing low-complex hard skills within the same occupational group is consistent with relative skill upgrading within occupations (see Proposition 3): skills associated with less automatable or complementary activities (soft skills) increase relative to skills concentrated in more automatable activities (low-complex hard skills), leading to an increase in the ratio $S^{soft}/S^{low-complex}$.

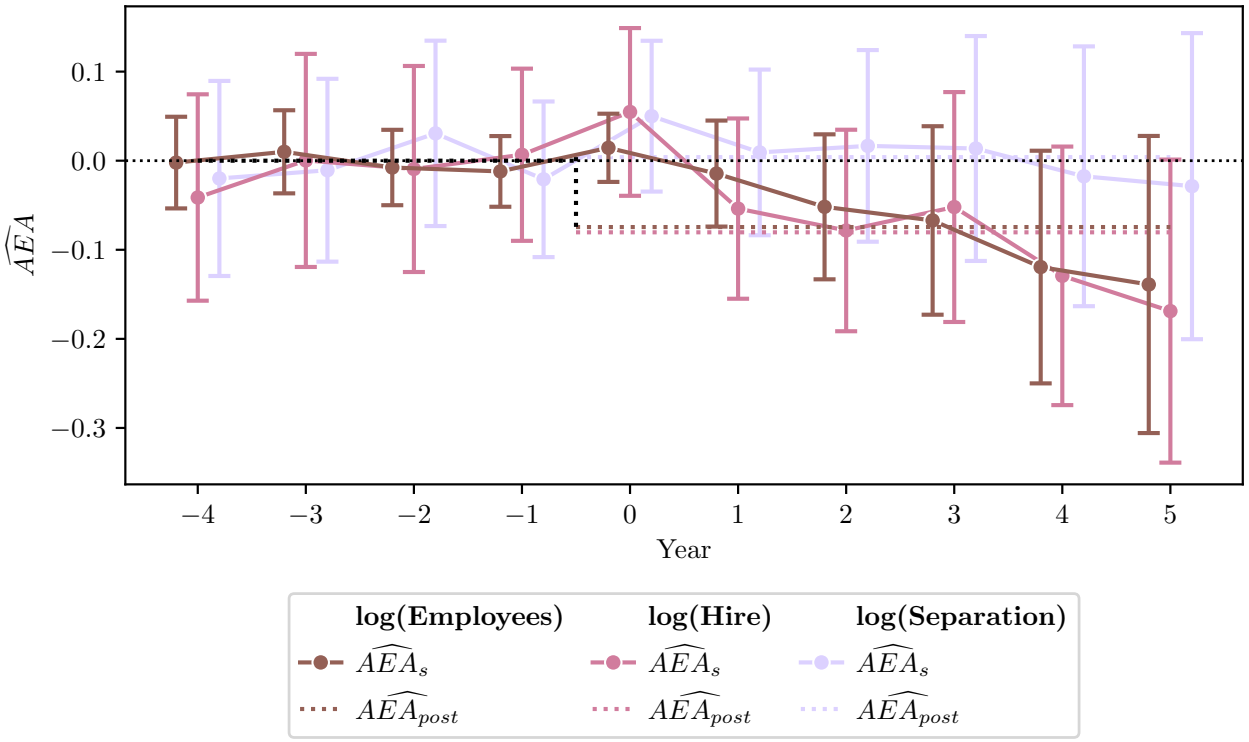
For administrative & service occupations, we find a large but statistically insignificant positive effect on low-complex hard skills. While this pattern is consistent with the descriptive decomposition, and suggests that hard-skill requirements may rise within surviving administrative and service roles, the estimate is imprecise, and the robustness checks in Section 4.3 show that this result is sensitive to sample variation.

For production & elementary occupations, all three estimates are close to zero and statistically insignificant, suggesting that skill requirements within these occupations remain broadly unchanged following automation.[35] In the context of our framework, the displacement mechanism described in Proposition 1, which predicts a decline in the within-occupation share of skills used intensively in tasks substituted by automation, does not appear to be a quantitatively important channel for the skills captured by our classification. Instead, adjustment in production & elementary occupations operate primarily through employment, as documented below, rather than through changes in skill requirements of new vacancies.

To provide evidence on the employment adjustments induced by automation, Figure 12 reports firm-level dynamic estimates from the Callaway and Sant'Anna (2021) approach applied to the same staggered event-study design, while Figure 13 presents the aggregate post-treatment estimates $\widehat{AEA}^{g}_{post}$ from equation (9) for employment measures at the broad occupational group level. Figure 12 shows that firms adopting automation reduce total employment primarily through lower hiring, while yearly separations remain at their pre-automation level. This pattern results in post-automation employment that is, on average,

[35] We cannot rule out that measurement error in vacancy-based skill measures for production & elementary occupations contributes to the imprecision of these estimates (see Section 3.3), but their small magnitudes suggest that any undetected effects are unlikely to be economically meaningful.

Figure 12: Dynamic Effects of Automation on Employment



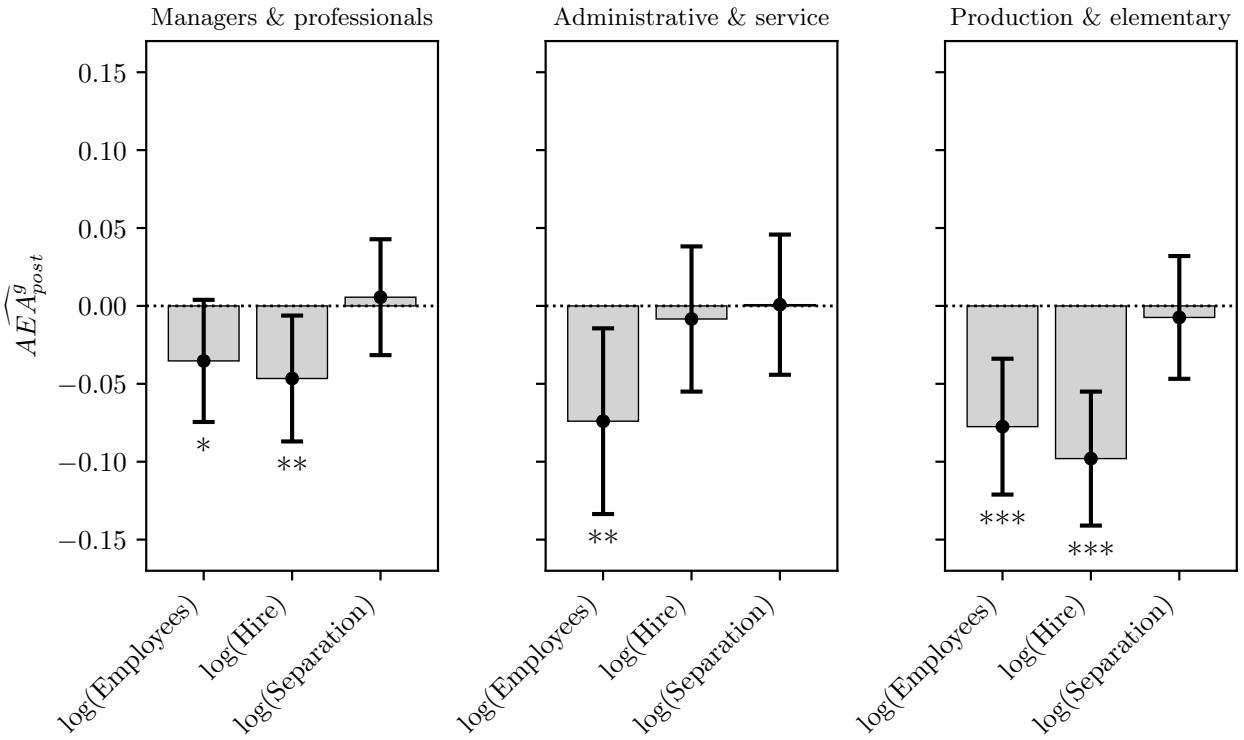
Notes: This figure reports dynamic event-study estimates of the effect of automation on firm-level employment, hiring, and separations, using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021). Event-time coefficients $\widehat{AEA}_s^g = \sum_c w_c \widehat{AEA}^g(c, c+s)$ are plotted for event time $s \in \{-4, \ldots, 5\}$, where $s = 0$ is the year of the automation event. The estimator uses not-yet-adopter firms as the comparison group. The outcome variables are the natural logarithm of the number of employees, the number of hires (workers not employed by the firm in the previous year), and the number of separations (workers no longer employed by the firm in the following year), all measured at the firm level. The dotted horizontal lines represent the corresponding aggregate post-treatment estimate $\widehat{AEA}_{post}^g$ from Figure 13. The lead coefficients ($s < 0$) provide evidence on the parallel trends and no-anticipation assumptions underlying the identification strategy. The sample consists of 20,317 firm-year observations from 1,847 automating firms. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level.

around 8% lower than its pre-automation level.[36] Figure 13 provides a more detailed view of how this contraction in total employment occurs. Automating firms reduce employment in production & elementary occupations through lower hiring, while the separation margin remains unchanged. The evidence for administrative & service workers suggests that employment declines without substantial adjustment through external labour market flows, as hires and separations show little change. By contrast, managers & professionals experience lower hiring but relatively stable employment. This pattern is consistent with some workers

[36] Aghion et al. (2022) survey the literature and contrast two views on the employment effects of automation: a displacement view, in which automation directly substitutes capital for labour, and a productivity view, in which automation-driven gains in productivity and market expansion can offset job losses within the automating firm. Our firm-level estimates capture the net effect of both channels on automating firms, without distinguishing between them, and point to a negative overall effect of automation on adopters.

in administrative occupations being reallocated internally into professional roles. The coefficients on leads in Figure 12 suggest parallel trends and no anticipation for adopters and not-yet-adopters. However, identifying assumptions underlying these estimates, including pre-treatment dynamics and robustness to controls, are discusses in more detail in Section 4.3.

Figure 13: Effects of Automation on Employment

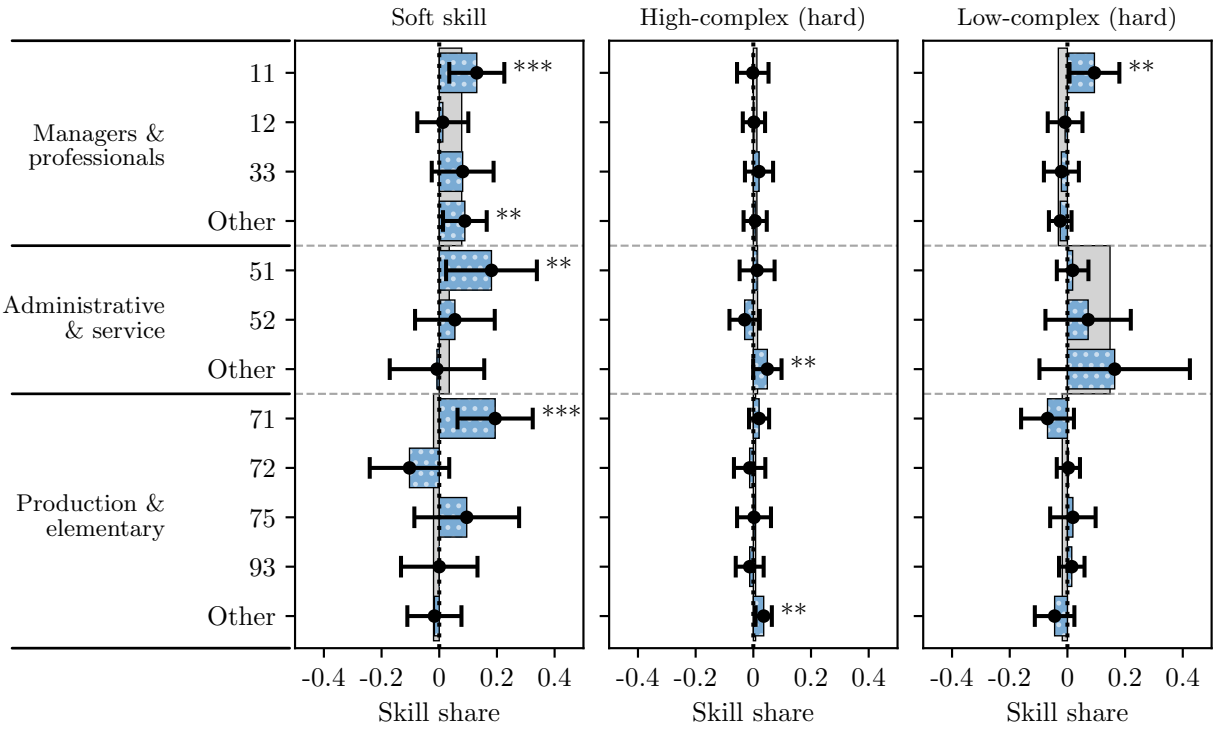


Notes: This figure reports the aggregate post-treatment Average Effect of Automation ($\widehat{AEA}^g_{post}$) on employment measures, estimated using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021) as described in Equation (9). The estimator uses not-yet-adopter firms as the comparison group. The outcome variables are the natural logarithm of the number of employees, the number of hires (workers not employed by the firm in the previous year), and the number of separations (workers no longer employed by the firm in the following year). Each panel displays estimates from a separate event study for one of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). See Figure 12 for the corresponding firm-level dynamic estimates. The sample consists of 20,317 firm-year observations from 1,847 automating firms. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Three main takeaways emerge from the estimated effects of automation on within-group changes in skill demand. First, automation does not have a uniform skill bias: its effects vary across skill types and occupational groups. Second, the substantial heterogeneity across occupational groups, already evident in the decomposition in Section 3.5, persists under our identification strategy. Figure A3 in the Appendix reports the Average Effect of Automation

on skill demand at the firm level, demonstrating that changes are masked at such level of aggregation. This evidence is consistent with point (iv) of the conceptual framework in Section 2: aggregate ΔS^k can be small even when occupation-specific ΔS^k_o terms are large, if they carry opposite signs across occupations. Third, within-occupation changes in skill demand play a central role primarily among occupations with tasks complementary to automation, such as managers & professionals, whereas compositional shifts dominate in the production & elementary group.

Figure 14: Within-Occupation Changes in Skill Demand by Detailed Occupation



Notes: This figure reports occupation-level estimates of the Average Effect of Automation ($\widehat{AEA}^o_{post}$) on within-occupation skill demand, using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021). Occupation-specific estimates (dotted blue bars) isolate genuine within-occupation changes in skill demand, abstracting from reallocation across occupations within the same group. The gray background bars reproduce the corresponding group-level estimates $\widehat{AEA}^g_{post}$ from Figure 11 for comparison. All outcomes are weighted by the number of vacancies. Within each occupational group, the most common two-digit occupations that jointly account for at least 50% of total vacancies in the group are shown individually; the remaining occupations are estimated jointly as the "Other" rows. The occupations shown are: 11 (Chief executives), 12 (Administrative managers), 33 (Business professionals), 51 (Personal service workers), 52 (Sales workers), 71 (Construction trade workers), 72 (Metal & machinery workers), 75 (Food & other craft trade workers). Observations vary by occupation: 2,618–18,656 firm-year observations from 238–1,696 firms; skill content based on 1,202–14,892 vacancies. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Detailed Occupations. The group-level estimates, $\widehat{AEA}^g_{post}$, reported in Figure 11 capture the effect of automation on skill demand within each broad occupational group, but

they do not distinguish whether this effect operates through genuine changes in the skill content of individual occupations or through the reallocation of vacancies across 2-digit occupations within the same group.[37] To examine whether automation induces genuine within-occupation changes in the composition of skill demand, and to identify which specific occupations experience the largest changes, Figure 14 presents results from the same staggered difference-in-differences approach, estimated separately at the level of individual two-digit ISCO occupations. The gray background bars reproduce the corresponding group-level $\widehat{AEA^g_{post}}$ estimates from Figure 11.

Among managers & professionals, the occupation-level estimates reveal that the group-level increase in soft-skill demand is not an artifact of reallocation across managerial occupations. Among the most common occupations, which jointly account for at least 50% of total vacancies in the group, Chief Executives (ISCO 11) and Business Professionals (ISCO 33) show sizable increases in within-occupation soft-skill requirements. The remaining occupations in the same group, grouped under the label "Other," exhibit the same significant pattern. This evidence corroborates the group-level finding: within executive and professional roles, where coordination, strategic oversight, and interpersonal engagement are core activities, automation shifts the skill mix most strongly toward soft skills. For this same group of occupations, the largely null changes observed in hard-skill demand, both high- and low-complexity, are also confirmed at the detailed two-digit level. In particular, automation does not affect within-role demand for high-complex hard skills such as engineering, finance, specialized software, and other advanced skills acquired through formal education.

Among administrative & service roles, the within-occupation change in soft skills is mostly null, in line with the group-level evidence in Figure 11. An exception is Personal Service Workers (ISCO 51), for whom automation induces a statistically significant increase

[37] For instance, a shift in the composition of vacancies in the managers & professionals group away from Business and Administration Professionals (ISCO 24) and toward Administrative Managers (ISCO 12) could generate a change in the group-level soft-skill share even if no individual occupation experienced any change in its skill requirements.

in soft-skill requirements, similarly to manager & professionals, suggesting that for such occupations the technology generates a reweighting toward complementary tasks for which interpersonal skills are salient. The change in hard-skill requirements within detailed administrative & service occupations is also broadly in line with the aggregate effect at group level and the descriptive decomposition in Section 3.5, showing some indication, although not statistically significant, that the hard-skill requirements of the remaining administrative roles increase.

Finally, the evidence from detailed roles in the production & elementary occupations cluster is also in line with that at group level, showing mostly null changes in skill requirements for new vacancies.[38]

4.2 Automation & Workforce Composition

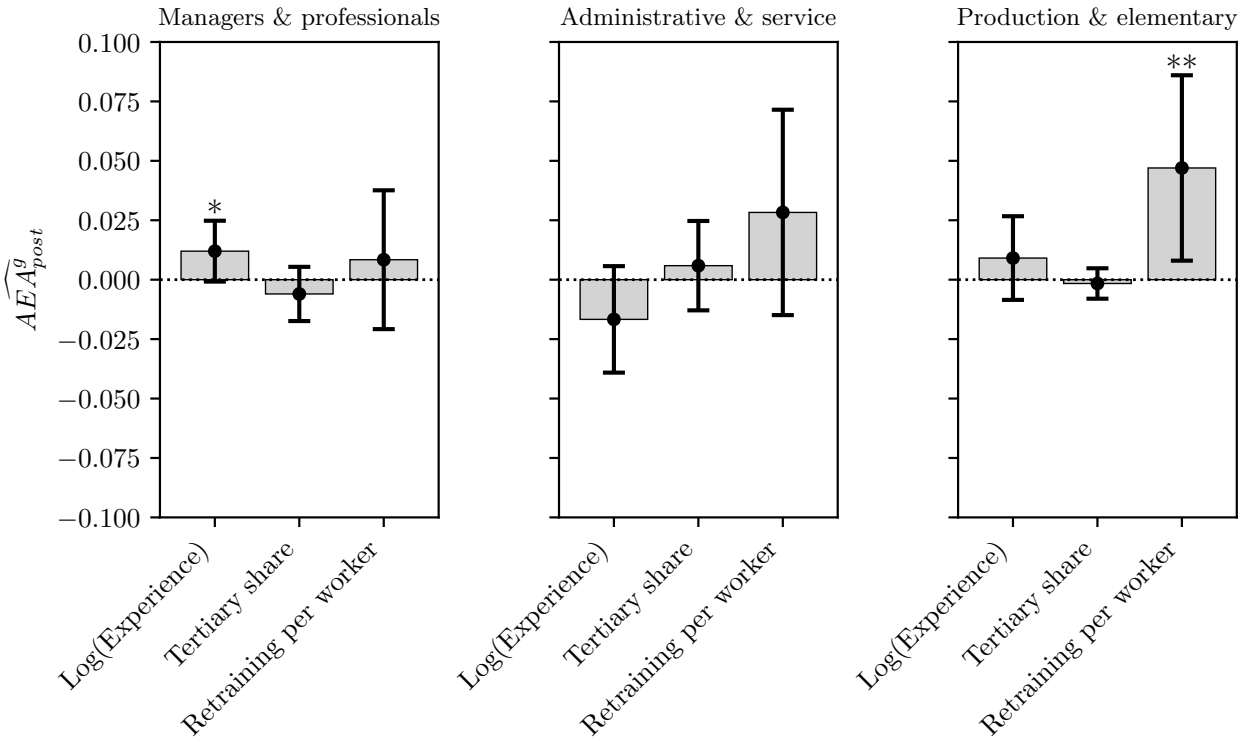
The previous section examined how automation affects the skills that firms seek in new hires, as revealed by job vacancy postings. However, vacancies capture desired skill requirements rather than realized workforce composition, and automation may also alter the nature of work for incumbent workers who remain in the firm. To provide evidence on how automation affects the actual skills and characteristics of the workforce, we apply the same empirical strategy presented in the beginning of Section 4 to complementary measures derived from administrative registers: average labour market experience, which serves as a proxy for the general-purpose competencies captured by soft skills in the vacancy data; the share of tertiary-educated workers, corresponding to the high-complex hard skills typically acquired through formal education; and retraining intensity, which captures skill adjustment among workers in roles where low-complex hard skills are prevalent.[39]

Figure 15 reports the Average Effect of Automation $\widehat{AEA^g_{post}}$ for each broad occupational

[38] Some heterogeneity is nevertheless present within the production cluster in terms of soft-skill requirements. Construction Trade Workers (ISCO 71), arguably a more non-routine manual occupation, become more interactive, while more routine manual roles, such as Metal, Machinery and Related Trades Workers (ISCO 72), show no such change.

[39] See Section 3.4 for details on these measures and their connection to the vacancy-based skill categories.

Figure 15: Effects of Automation on Workforce Composition by Occupational Group



Notes: This figure reports the aggregate post-treatment Average Effect of Automation ($\widehat{AEA}^{g}_{post}$) on register-based measures of workforce composition, estimated using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021) as described in Equation (9). The estimator uses not-yet-adopter firms as the comparison group. The outcome variables are: the natural logarithm of average years of labour market experience, the share of workers with any tertiary-level degree, and the number of retraining courses taken per worker. Each panel displays estimates from a separate event study for one of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). The sample consists of 20,317 firm-year observations from 1,847 automating firms. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

group. The left-hand panel suggests that average labour market experience increases by 1% among managers & professionals (significant at 10%). This is consistent with the finding that automation increases demand for soft skills in these occupations, not by raising demand for formal competencies, but through greater reliance on workers with more accumulated experience.

The absence of within-occupation changes in skill demand for administrative & service as well as production & elementary occupations, discussed in the previous section, is consistent with the contraction of these occupations shown in Figure 13. At the same time, automation may still alter the tasks performed by incumbent workers who remain in the firm. The middle and right-hand panels of Figure 15 show that automation does not significantly alter

the educational level of the workforce, but increases retraining intensity among production & elementary workers by 5 percentage points relative to the pre-automation period. This suggests that automation changes the nature of work in these occupations, prompting firms to adjust the skills of production workers through targeted retraining. The most common courses among production & elementary workers, reported in Table 2, are consistent with this interpretation: they include courses in transportation, metalwork, and welding, reflecting the competencies required to operate alongside new machinery.

4.3 Robustness & additional checks

Automation events. Automation events are intended to capture firms' tendency to lump automation investments (Bessen et al., 2023). For empirical implementation, this requires a threshold-based definition. In this paper, we define an automation event as an instance in which a firm's largest machinery investment exceeds four times its average investment over the panel. We assess the robustness of this definition by considering alternative thresholds of three and five. As illustrated by the distribution of relative investments in Figure A1 in the Appendix, changing the cutoff by one alters the sample of automating firms by about 20%. The results of these robustness checks, together with the main specification, are reported in Figure A7 in the Appendix. Overall, the results for the alternative cutoffs are similar to the main findings. Most importantly, the positive and statistically significant coefficient for managers & professionals remains robust across all cutoff values. By contrast, the coefficient on low-complex hard skills for managers & professionals, which is significant at the 10% level in the baseline specification, is no longer significant under the alternative cutoffs.

Parallel trends & no anticipation. The staggered difference-in-differences approach relies on the assumptions of parallel trends between adopters and not-yet-adopters prior to automation, and of no anticipation of automation affecting outcomes before the event occurs. Because automation is a firm-level choice, a concern is that firms may begin adjusting hiring or workforce composition in anticipation of future technology adoption. In that case,

estimated post-treatment effects could partly capture anticipatory responses rather than the causal effect of automation. To assess this possibility, we examine pre-treatment dynamics in the Callaway and Sant'Anna (2021) event-study framework. Figure A4 reports the dynamic estimates for the three skill types across the three occupational groups over the event window $[-4, +5]$. Inspection of the lead coefficients reveals no systematic pre-treatment divergence between adopters and not-yet-adopters. Similar reassuring evidence is found for total employment, hires, and separations in Figure 12. For all other dependent variables, Table A7 reports p-values for Wald χ^2 tests of the null hypothesis of no pre-trends across the four pre-event periods $\tau \in \{-4, -3, -2, -1\}$: none of the joint tests, nor any of the individual period coefficients, is statistically significant. This evidence alleviates concerns that our estimates are driven by pre-existing trends or anticipatory firm behavior.

Control variables. The absence of pre-trends suggests that the baseline approach relies on an appropriate comparison group that follows similar dynamics prior to automation. As a robustness check, we re-estimate the main specifications including pre-automation covariates: industry at the 1-digit NACE level, quartiles of firm employment, and quartiles of the provincial share of tertiary-educated workers, capturing differences in sectoral composition, firm size, and local labour supply conditions. As shown in Figure A5 in the Appendix, the resulting estimates remain very similar to the baseline results, mitigating concerns that the estimates are driven by differential outcome dynamics across firms with different observable characteristics.

Dynamic estimates and statistical power. Figure A4 in the Appendix reports dynamic event-study estimates of the main specification. None of the event-time-specific coefficients $\widehat{AEA_s^g}$ reaches statistical significance individually. In Section 2 of the Online Appendix, we show that this is consistent with limited statistical power at the event-time level: a post-hoc power calculation yields a power of only 17% for the event-time estimate of soft skills among managers & professionals, compared to 82.5% for the corresponding aggregate post-treatment estimate $\widehat{AEA_{post}^g}$. Reassuringly, the $\widehat{AEA_s^g}$ estimates align in magnitude

with $\widehat{AEA^g_{post}}$ and suggest a gradual shift in skill demand, with the event-time coefficients reaching the size of the aggregate estimate by $\tau = 2$.

Additional checks. In Section 3 of the Online Appendix, we examine the results for individual skill categories following Deming and Kahn (2017). The increased demand for soft skills among managers & professionals appears to be driven primarily by social and noncognitive skills. There are also some indications of upskilling and more digitalized work among administrative & service workers.

In Section 4 of the Online Appendix, we investigate the results using an alternative, import-based definition of automation events. The subset of firms that automate through direct import and installation of technologies shows somewhat different patterns, with no significant increase in the demand for soft skills among managers & professionals.

As a further check, we verify that our focus on relative skill shares $S^k_{o,t}$ does not capture important alternative channels of adjustment. Figure A6 in the Appendix reports event-study estimates with the number of vacancies and the number of skills per vacancy as outcomes and generally shows null effects. The only significant coefficient is an increase in the number of skills per vacancy among administrative & service occupations. This result, combined with the large but imprecise coefficients in the main results of Figure 11, suggests that the effect of automation on this occupational group remains less well identified than for the other groups.

Overall, the positive effect of automation on the demand for soft skills among managers & professionals is robust across all specifications, with the exception of import-based automation events. The decline in low-complex hard skills for the same group, significant at the 10% level in the baseline, is sensitive to alternative specifications and should be interpreted with caution.

5 Conclusion

This paper examines how firms adjust skill demand in response to automation. Combining firm-level automation data with job vacancy postings and worker-level administrative records from Denmark, we document that the effects of automation on skill demand are systematic but vary markedly across occupations. A descriptive decomposition shows that within-occupation changes account for 14–39% of total skill demand adjustment, establishing that automation operates not only through occupational reallocation but also by reshaping the skill content of existing jobs. Using a staggered difference-in-differences design, we estimate the causal effects of automation on multidimensional skill demand across the occupational hierarchy. For managers & professionals, automation increases the demand for soft skills, shifting the within-occupation skill mix toward interpersonal and cognitive competencies. For production & elementary workers, by contrast, we find no significant within-occupation changes in vacancy skill requirements; instead, adjustment operates primarily through reduced hiring and the resulting contraction in employment. These patterns point to distinct margins of adjustment. In higher-tier occupations, automation is accompanied by changes in the composition of skill requirements within jobs, consistent with a shift toward activities requiring coordination, interaction, and problem-solving. In lower-tier occupations, by contrast, adjustment occurs primarily through reductions in hiring, while the skill content of new vacancies remains largely unchanged. At the same time, increased retraining among incumbent production workers suggests that firms respond to technological change by updating worker capabilities even when the formal skill requirements of new positions change little.

To interpret these findings, we draw on a task-based conceptual framework that extends canonical models of automation to multidimensional skill demand within occupations. The framework distinguishes between displacement of skills used in automatable tasks and reweighting toward skills used in complementary activities, providing a structured lens through which

to read the empirical patterns. The increased demand for soft skills among managers & professionals is consistent with automation raising the importance of coordination and oversight tasks that complement automated processes, even when few tasks in these occupations are directly automated. The null effects on vacancy skill requirements for production workers, alongside significant employment contraction and increased retraining, suggest that in these occupations automation operates primarily through the extensive margin rather than through within-occupation task recomposition.

Our results carry implications for how policymakers and firms respond to technological change. The finding that adjustment often takes place within occupations, or through changes in hiring and retraining rather than in posted skill requirements, means that standard measures based on occupational employment shares may not fully capture how automation reshapes work. For higher-tier occupations, our evidence that automation raises demand for interpersonal and cognitive skills, and that firms respond by relying on more experienced workers, resonates with recent work showing that task removal can raise expertise requirements within occupations (Autor and Thompson, 2025). This suggests that experience-based skill development and the accumulation of judgment and coordination capabilities may be as important as formal technical training in preparing workers for automation. For lower-tier occupations, the significant increase in retraining highlights the value of accessible, occupation-specific reskilling programs that equip incumbent workers with the competencies required alongside new technologies. Denmark's extensive system of adult education, which facilitates targeted retraining for both employed and unemployed workers, provides a model for how such programs can support within-occupation adaptation.

More broadly, the evidence in this paper suggests that the impact of automation depends on the task structure of occupations and operates through multiple margins: employment, task composition, and skill demand. An important avenue for future research is to examine how these within-occupation skill adjustments translate into longer-term outcomes for workers, including wage dynamics, career progression, and the returns to different types of skills.

Understanding how the reconfiguration of work documented here affects individual workers over time is central to assessing the full consequences of technological change and to designing policies that support effective adaptation.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.

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

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.

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

Acemoglu, D. and Restrepo, P. (2021). Demographics and Automation. *The Review of Economic Studies*, 89(1):1–44.

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

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2022). The direct and indirect effects of automation on employment: A survey of the recent literature. In Ing, L. Y. and Grossman, G. M., editors, *Robots and AI: A New Economic Era*. Routledge, Oxon and New York.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2023). Modern manufacturing capital, labor demand and product market dynamics: evidence from France. CEP Discussion Papers dp1910, Centre for Economic Performance, LSE.

Alekseeva, L., Azar, J., Giné, M., Samila, S., and Taska, B. (2021). The demand for ai skills in the labor market. *Labour Economics*, 71:102002.

Atalay, E., Phongthiengtham, P., Sotelo, S., and Tannenbaum, D. (2020). The evolution of work in the united states. *American Economic Journal: Applied Economics*, 12(2):1–34.

Autor, D. and Thompson, N. (2025). Expertise. *Journal of the European Economic Association*, 23(4):1203–1271.

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

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

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing inequality: Have computers changed the labor market? *The Quarterly Journal of Economics*, 113(4):1169–1213.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The skill content of recent technological change: An empirical exploration. *The Quarterly journal of economics*, 118(4):1279–1333.

Battisti, M., Dustmann, C., and Schönberg, U. (2023). Technological and organizational change and the careers of workers. *Journal of the European Economic Association*, 21(4):1551–1594.

Beaudry, P., Green, D. A., and Sand, B. M. (2016). The great reversal in the demand for skill and cognitive tasks. *Journal of Labor Economics*, 34(S1):S199–S247.

Bessen, J., Goos, M., Salomons, A., and Van den Berge, W. (2023). What happens to workers at firms that automate? *Review of Economics and Statistics*, pages 1–45.

Bonfiglioli, A., Crinò, R., Fadinger, H., and Gancia, G. (2024). Robot imports and firm-level outcomes. *The Economic Journal*, 134(664):3428–3444.

Brynjolfsson, E., Buffington, C., Goldschlag, N., Li, J. F., Miranda, J., and Seamans, R. (2023). The characteristics and geographic distribution of robot hubs in u.s. manufacturing establishments. Working Paper 31062, National Bureau of Economic Research.

Békés, G., Bonfiglioli, A., Crinò, R., and Gancia, G. (2025). Integrators and robot adoption: Facts from hungary. Discussion Paper 20026, CEPR.

Caines, C., Hoffmann, F., and Kambourov, G. (2017). Complex-task biased technological change and the labor market. *Review of Economic Dynamics*, 25:298–319.

Callaway, B. and Sant'Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230. Themed Issue: Treatment Effect 1.

de Souza, G. and Li, H. (2023). Robots, tools, and jobs: Evidence from brazilian labor markets. Working Paper 10813, CESifo.

Deming, D. and Kahn, L. B. (2017). Skill requirements across firms and labor markets: Evidence from job postings for professionals. *Journal of Labor Economics*, 36(S1):S337–S369.

Deming, D. and Silliman, M. (2025). Chapter 2 - skills and human capital in the labor market. In Dustmann, C. and Lemieux, T., editors, *Handbook of Labor Economics*, volume 6 of *Handbook of Labor Economics*, pages 115–157. Elsevier.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640.

Deming, D. J. and Noray, K. (2020). Earnings dynamics, changing job skills, and stem careers. *The Quarterly Journal of Economics*, 135(4):1965–2005.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding.

Domini, G., Grazzi, M., Moschella, D., and Treibich, T. (2021). Threats and opportunities in the digital era: Automation spikes and employment dynamics. *Research Policy*, 50(7):104137.

Edin, P.-A., Fredriksson, P., Nybom, M., and Öckert, B. (2022). The rising return to noncognitive skill. *American Economic Journal: Applied Economics*, 14(2):78–100.

Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277. Themed Issue: Treatment Effect 1.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American Economic Review*, 104(8):2509–2526.

Graetz, G. and Michaels, G. (2018). Robots at work. *Review of Economics and Statistics*, 100(5):753–768.

Grönqvist, E. and Lindqvist, E. (2016). The making of a manager: Evidence from military officer training. *Journal of Labor Economics*, 34(4):869–898.

Heckman, J. J. (1995). Lessons from the bell curve. *Journal of Political Economy*, 103(5):1091–1120.

Heckman, J. J. and Kautz, T. (2012). Hard evidence on soft skills. *Labour Economics*, 19(4):451–464.

Hershbein, B. and Kahn, L. B. (2018). Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–72.

Hirvonen, J., Stenhammar, A., and Tuhkuri, J. (2025). What do technology grants do? Working paper, mimeo.

Humlum, A. (2022). Robot adoption and labor market dynamics. Technical report, Rockwool Foundation Research Unit Berlin, Germany.

Izadi, R. and Tuhkuri, J. (2024). Evolving returns to personality. Technical Report 16/24, Rockwool Foundation Berlin (RF Berlin) and Centre for Research & Analysis of Migration (CReAM), Berlin and London.

Lindqvist, E. and Vestman, R. (2011). The labor market returns to cognitive and noncogni-

tive ability: Evidence from the swedish enlistment. *American Economic Journal: Applied Economics*, 3(1):101–28.

Petrova, M., Schubert, G., Taska, B., and Yildirim, P. (2024). Automation, career values, and political preferences. Working Paper 32655, National Bureau of Economic Research.

Reimers, N. and Gurevych, I. (2019). Sentence-bert: Sentence embeddings using siamese bert-networks.

Reimers, N. and Gurevych, I. (2020). Making monolingual sentence embeddings multilingual using knowledge distillation. In Webber, B., Cohn, T., He, Y., and Liu, Y., editors, *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4512–4525, Online. Association for Computational Linguistics.

Salomons, A., vom Baur, C., and Zierahn-Weilage, U. (2025). Expertise at work: New technologies, new skills, and worker impacts. *IZA Discussion Paper*, (18248).

Spitz-Oener, A. (2006). Technical change, job tasks, and rising educational demands: Looking outside the wage structure. *Journal of Labor Economics*, 24(2):235–270.

Appendix

A Tables & Figures

Table A1
Automation Investment by Industry

		No. Firms	No. Automating firms	Avg. Investments
C	Manufacturing	2,880	1,348	2,423.0
F	Construction	2,266	696	331.0
G	Wholesale and retail trade	5,187	1,680	245.0
H	Transportation and storage	668	351	6,353.5
I	Accommodation and food service activities	1,337	116	17.5
J	Publishing, broadcasting & content	1,071	206	2,726.0
M	Real estate activities	1,829	363	175.6
N	Professional, scientific and technical activities	986	195	2,477.3
S	Arts, sports and recreation	65	27	84.3

Notes: This table reports the distribution of automation activity across industries classified by NACE letter codes. For each industry, the table shows the total number of firms, the number of firms experiencing an automation event as defined in Equation (7), and the average level of machinery investment (in thousands of DKK). The sample includes nine sectors in which machinery investment is observed and plausibly reflects automation technologies rather than infrastructure-related capital; see Table A2 for the full list of included and excluded sectors and the criteria for exclusion. The variation in average investment across industries reflects differences in the capital intensity of production processes across sectors.

TABLE A2
Included Sectors and Exclusion Criteria

	Sector	Included	Excluded		
			No machinery	Financial	Infrastructure
A	Agriculture, Forestry and Fishing		✓		
B	Mining and Quarrying				✓
C	Manufacturing	✓			
D	Electricity, Gas, Steam and Air Conditioning Supply				✓
E	Water Supply; Sewerage, Waste Management				✓
F	Construction	✓			
G	Wholesale and Retail Trade; Repair of Motor Vehicles	✓			
H	Transportation and Storage	✓			
I	Accommodation and Food Service Activities	✓			
J	Publishing, Broadcasting, and Content Production and Distribution	✓			
K	Information Service Activities		✓		
L	Financial and Insurance Activities			✓	
M	Real Estate Activities	✓			
N	Professional, Scientific and Technical Activities	✓			
O	Administrative and Support Service Activities		✓		
P	Public Administration and Defence; Compulsory Social Security		✓		
Q	Education		✓		
R	Human Health and Social Work Activities		✓		
S	Arts, Sports and Recreation	✓			
T	Other Service Activities		✓		
U	Activities of Households as Employers		✓		
V	Activities of Extraterritorial Organisations and Bodies		✓		

Notes: This table lists all NACE sectors and indicates whether each is included in or excluded from the analysis. Included sectors are those in which machinery investment is observed and is plausibly related to automation rather than infrastructure. Exclusion criteria fall into three categories: "No machinery" denotes industries with zero recorded machinery investments over the sample period; "Financial" denotes the financial and insurance sector, excluded due to the distinct nature of capital in these industries; "Infrastructure" denotes sectors where machinery investments are more likely to reflect infrastructure-related capital (e.g., pumps, heating systems) rather than automation equipment.

TABLE A3
Descriptive Statistics for Automating Firms

	Firm level		Managers & professionals		Administrative & service		Production & elementary	
	Mean	Std.	Mean	Std.	Mean	Std.	Mean	Std.
Number of vacancies	1.7	9.1	1.0	5.9	0.3	3.8	0.4	1.9
Number of skills	8.8	24.6	8.3	18.8	5.7	19.3	4.0	7.6
Skills per vacancy	0.70	1.04	0.53	1.00	0.20	0.65	0.25	0.67
Soft skill share	0.18	0.29	0.12	0.23	0.05	0.17	0.09	0.24
High-complex (hard) share	0.07	0.15	0.06	0.13	0.02	0.08	0.02	0.09
Low-complex (hard) share	0.10	0.18	0.07	0.15	0.03	0.11	0.03	0.12
Retraining	0.22	0.53	0.13	0.45	0.17	0.66	0.31	0.69
Tertiary	0.24	0.27	0.40	0.33	0.20	0.29	0.07	0.17
Experience	22.4	5.9	24.6	6.4	23.2	7.9	22.1	7.3
Employment	33.3	162.2	17.2	76.7	10.9	134.2	15.8	41.4
Share of employment			0.38	0.33	0.17	0.24	0.38	0.35
Hourly wage	237.7	93.8	293.5	131.3	203.9	58.1	197.8	41.0
Hires	4.0	16.1	2.2	11.5	1.2	5.8	1.9	6.8
Fires	3.5	13.3	1.9	8.5	1.1	6.2	1.6	5.5
Revenue	86,967	218,823						

Notes: This table reports descriptive statistics for the sample of automating firms used in the main analysis. All values are computed over the full firm-year panel (2008–2021). Statistics are reported at the firm level and separately for each of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). Vacancy-based variables (number of vacancies, number of skills, skills per vacancy, and skill shares) are derived from postings matched to automating firms. Skill shares report the share of each broad skill type in total skills demanded. Register-based variables include retraining, tertiary education share, experience, employment, hourly wage (in DKK), hires, separations, and revenue (in thousands of DKK). The share of employment reports the mean occupational group employment share relative to total firm employment.

TABLE A4
Skill Categories

Skill	Notes	Aggregated category
Cognitive		Soft skill
Social		Soft skill
Noncognitive		Soft skill
Management		Soft skill
Creativity		Soft skill
Project management		High-complex hard skill
Finance		High-complex hard skill
Specialized software	Combination of specialized software and AI + added corpus of programming skills	High-complex hard skill
Data	Combination of database, data analysis, business analysis + added corpus of all skills mentioning data not elsewhere classified	High-complex hard skill
Engineer		High-complex hard skill
Products		High-complex hard skill
Admin support	Combination of administrative support and writing	Low-complex hard skill
Customer support		Low-complex hard skill
Computer	Combination of computer, general software, and business systems	Low-complex hard skill
Tech support		Low-complex hard skill

Notes: This table lists the 15 skill categories used in the analysis, along with their aggregation into three broad skill types: soft skills, high-complex hard skills, and low-complex hard skills. The categories are based on the taxonomies of Deming and Noray (2020) and Deming and Kahn (2017). Some categories are constructed by merging highly correlated categories from the original Deming and Noray (2020) classification, as indicated in the "Notes" column. Skills are extracted from job vacancy texts using a sentence-embedding approach that measures semantic similarity between extracted phrases and the keyword corpus associated with each category.

TABLE A5
Top-10 Skill Words per Category

Raw skill	Translation	N. vacs		Deming and Kahn (2017)		Skill group
samarbejde	cooperation	390,202	→			
samarbejdsevner	collaboration skills	243,749	→	Social	→	
kommunikere	communicate	196,995	→			
selvstændigt	independently	628,127	→			
fleksibel	flexible	521,423	→			
struktureret	structured	436,630	→			
engageret	engaged	400,431	→	Noncognitive	→	Soft skill
selvstændig	independent	262,681	→			
positiv	positive	231,443	→			
ansvarsbevidst	responsible	197,230	→			
projektleder	project manager	100,024	→			
projektledelse	project management	74,921	→	Project management	→	
project management	project management	41,175	→			
kvalitetssikring	quality assurance	39,101	→			
økonomi	economy	76,054	→	Finance	→	High-complex (Hard skill)
regnskab	accounting	38,992	→			
ingeniør	engineer	54,988	→	Engineer	→	
engineering	engineering	41,514	→			
marketing	marketing	53,436	→	Products	→	
indkøb	purchase	32,851	→			
administration	administration	68,198	→	Admin support	→	
administrative opgaver	administrative tasks	50,527	→			
serviceminded	service minded	310,361	→			
service	service	124,435	→			
kundeservice	customer service	91,645	→	Customer support	→	Low-complex (Hard skill)
sælge	sell	49,040	→			
serviceorienteret	service oriented	37,433	→			
it	it	247,511	→			
excel	excel	133,374	→	Computer	→	
sap	sap	66,136	→			

Notes: This table displays the ten most frequently classified skill words for a subset of skill categories, as they appear in the raw texts of job vacancies. For each word, the table reports the original Danish term, its English translation, and the number of vacancies in which the word appears. Skill words are assigned to categories using the sentence-embedding classification procedure described in Section 3.3. The mapping from individual skill categories to the broad groupings of Deming and Kahn (2017)—soft skills, high-complex hard skills, and low-complex hard skills—is indicated in the rightmost columns.

TABLE A6
Examples of Production-Related Skills

Skill category	Manufacturing terms
Soft skills	product creation, manage production, product management, ordering products, storage, item count, production planning, product replenishment, planning the production, manufacturing planning, ordering goods, efficient production
High-complex hard skill	craftsmanship, production engineering, industrial laboratory technician, industrial hoses, procurement production, industrial construction, craftsman management, manufacturing engineering, artisanal, production engineer, industrial tasks, merchandising, tradesman services, industrial electrician, production cost, industrial customers, tool maker, procurement for production, smart production, industrial engineering, production data, engineer in production, manufacturing data, purchase of production equipment, engineering direction, data factory, product engineering, smart factory, project-oriented production, production accounting, plant engineering, product order
Low-complex hard skill	service the production, service production equipment, serve products, servicing of production equipment, manufacturing support, production support, servicing

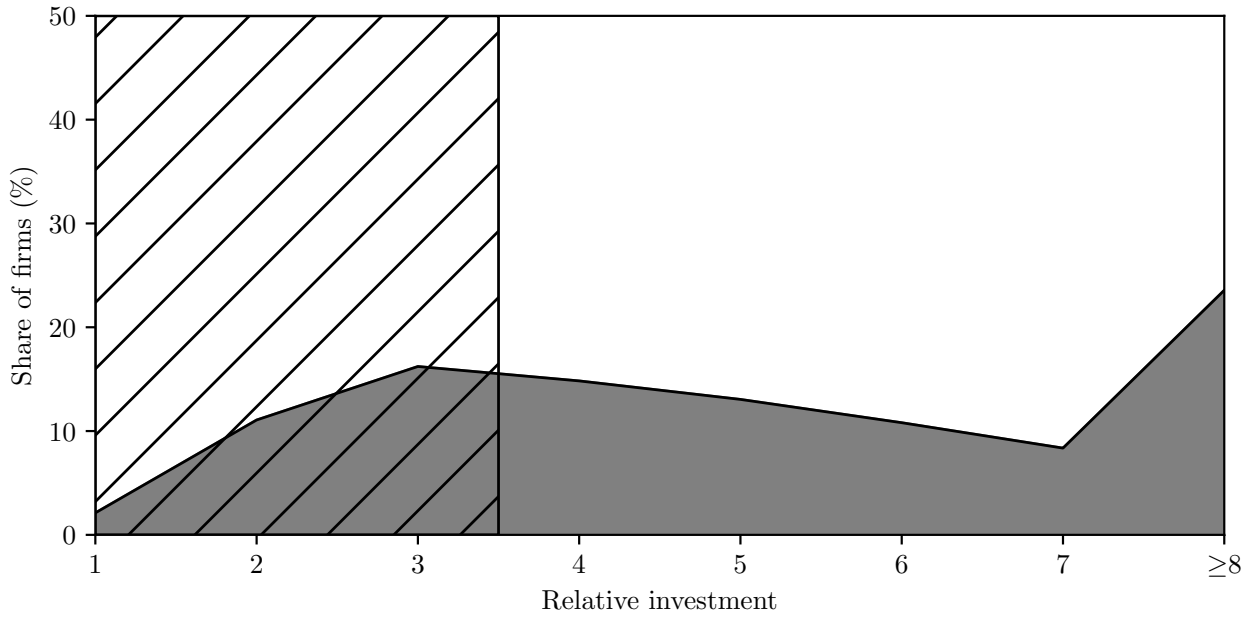
Notes: This table displays the 50 raw skill words (translated to English) with the highest semantic similarity to the terms "machine," "manufacturing," and "production," restricted to words appearing in at least 10 job vacancies. Unrelated words (e.g., "web production") are manually filtered out. The table illustrates the types of production-relevant skill content captured by each broad skill category. Among soft skills, production-related terms relate primarily to production management and planning. High-complex hard skills contain the richest set of production terms, spanning specific production methods and craftsmanship. Low-complex hard skills contain fewer such terms, mostly relating to machinery servicing and maintenance.

TABLE A7
Pre-Trends in Outcome Variables

Variable	P-value of Wald test (sign of mean coefficient)			
	Firm level	Managers & professionals	Administrative & service	Production & elementary
Soft skill	0.984 (+)	0.999 (-)	0.944 (+)	1.0 (+)
High complex (Hard)	0.961 (+)	0.961 (+)	0.932 (-)	0.885 (+)
Low complex (Hard)	0.992 (-)	0.972 (-)	0.994 (+)	0.917 (-)
Log(Experience)	0.973 (+)	0.965 (+)	0.947 (-)	0.996 (+)
Tertiary share	0.798 (-)	0.808 (-)	0.917 (-)	0.981 (+)
Retraining per worker	0.969 (-)	0.992 (-)	0.986 (-)	0.969 (+)
log(Total skills)	0.967 (-)	0.921 (-)	1.0 (+)	0.953 (-)
Total skills per vacancy	0.851 (-)	0.9 (+)	0.957 (+)	0.932 (+)
log(Vacancies)	0.979 (-)	0.914 (-)	0.944 (+)	0.91 (-)
log(Employees)	0.953 (-)	0.999 (-)	0.981 (+)	0.965 (-)
Employment share		0.964 (+)	0.996 (-)	0.988 (+)
log(Hire)	0.968 (-)	0.914 (-)	0.998 (-)	1.0 (-)
log(Fire)	0.945 (-)	0.993 (-)	0.998 (-)	0.937 (-)
log(revenue)	0.987 (-)			
log(Hourly wage)	0.987 (-)	0.903 (+)	0.972 (-)	0.982 (-)

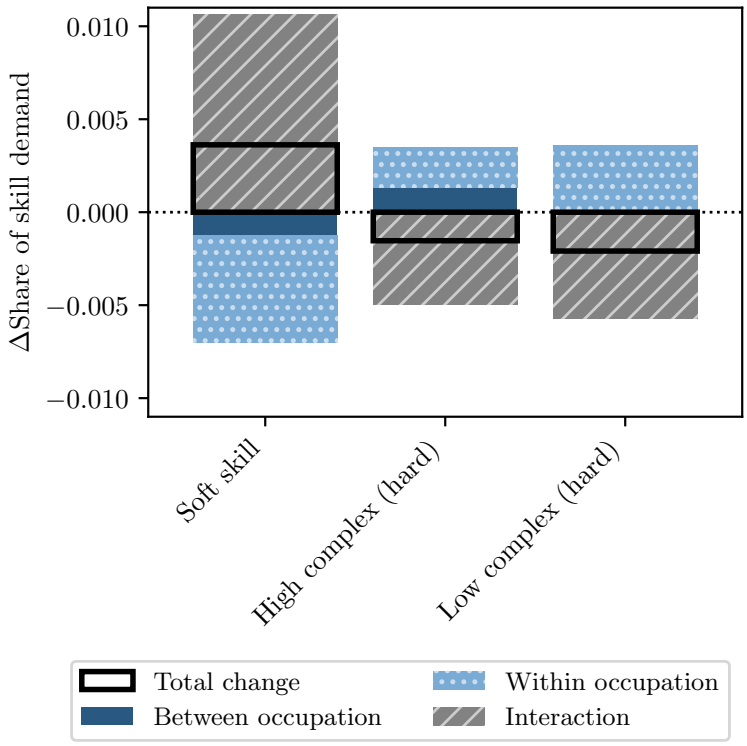
Notes: This table reports tests for pre-treatment trends in all dependent variables used in the main analysis. Each cell shows the p-value from a Wald $\chi^2(4)$ test of the null hypothesis that all four pre-event coefficients are jointly zero over $\tau \in \{-4, -3, -2, -1\}$, based on the dynamic event-study estimates from the Callaway and Sant'Anna (2021) staggered difference-in-differences estimator using not-yet-adopter firms as the comparison group. The sign in parentheses indicates the direction of the mean pre-event coefficient. Tests are reported at the firm level and separately for each occupational group (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). No joint test, nor any individual period coefficient, is statistically significant at conventional levels, supporting the parallel trends and no-anticipation assumptions underlying the identification strategy. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure A1: Distribution of Relative Investments per Worker



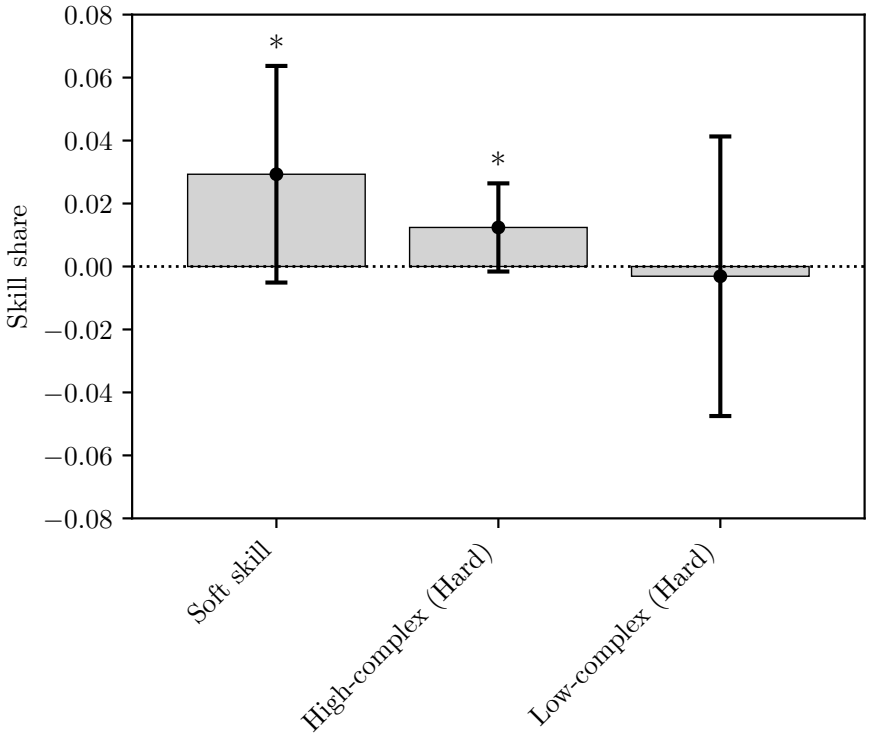
Notes: This figure displays the distribution of each firm's largest machinery investment over the sample period, expressed as a ratio to the firm's average investment. Investments below 15,000 EUR (112,000 DKK) are excluded. The cumulative area to the right of a given relative investment value represents the share of firm-years that are assigned automation events for a given automation event definition. The hatched area represents firms that are not considered as having an automation event in the main analysis.

Figure A2: Firm-Level Decomposition of Skill Demand



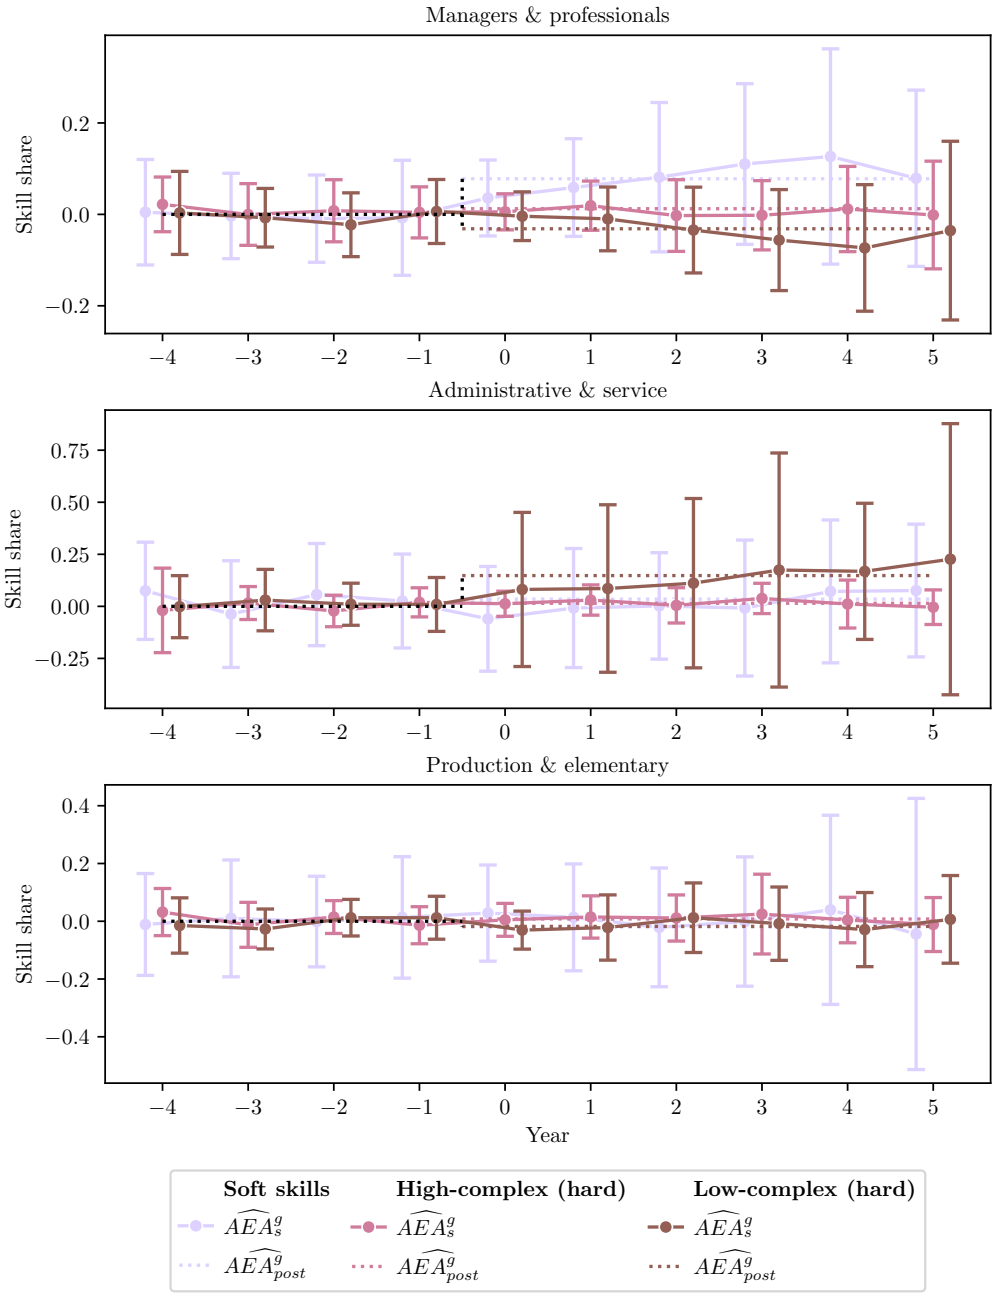
Notes: This figure displays the firm-level decomposition of changes in skill demand for automating firms, following Equation (6). Each bar reports the total change in the share of skill demand (ΔS^k) along with its three components: the between-occupation component ($\sum_o \Delta w_o \, S^k_{o,t_0}$), the within-occupation component ($\sum_o w_{o,t_0} \, \Delta S^k_o$), and the interaction component ($\sum_o \Delta S^k_o \, \Delta w_o$). Skill shares $S^k_{o,t}$ are measured at the 2-digit ISCO-08 occupation level; employment shares $w_{o,t}$ are measured as the share of vacancies posted for occupation o. Pre- and post-automation periods are defined by pooling vacancies over $\tau \in \{-3, -2, -1\}$ and $\tau \in \{1, 2, 3\}$, respectively. The change in share of vacancies is reported in percentage points. This figure aggregates over all occupations without distinguishing between occupational groups.

Figure A3: Firm-Level Effects of Automation on Skill Demand



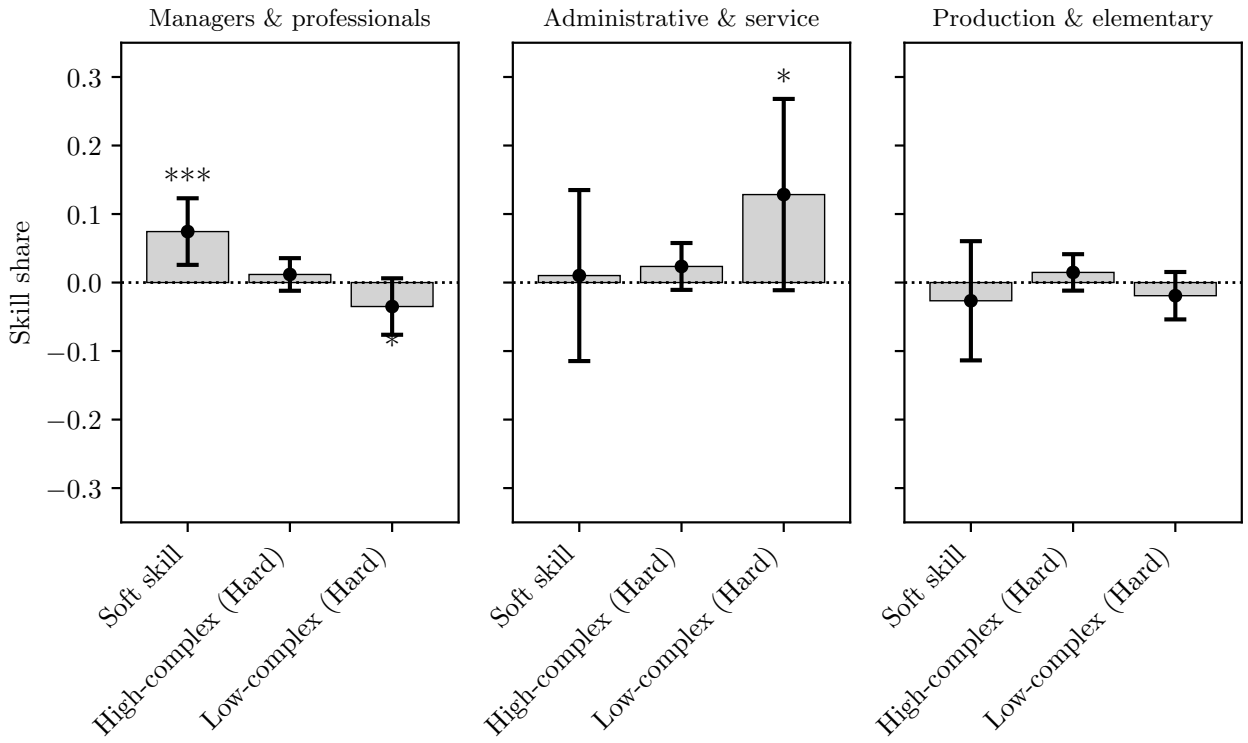
Notes: This figure reports the aggregate post-treatment Average Effect of Automation ($\widehat{AEA}_{post}$) on skill demand at the firm level, estimated using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021) as described in Equation (9). The estimator uses not-yet-adopter firms as the comparison group. The outcome variable is the share of each broad skill type (soft skills, high-complex hard skills, low-complex hard skills) in total skills demanded at the firm level. All outcomes are weighted by the number of vacancies. The sample consists of 14,051 firm-year observations from 3,571 automating firms, based on 58,728 vacancies. Full coefficient table reported in Table 2 of the Online Appendix. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure A4: Dynamic Effects of Automation on Skill Demand by Occupational Group



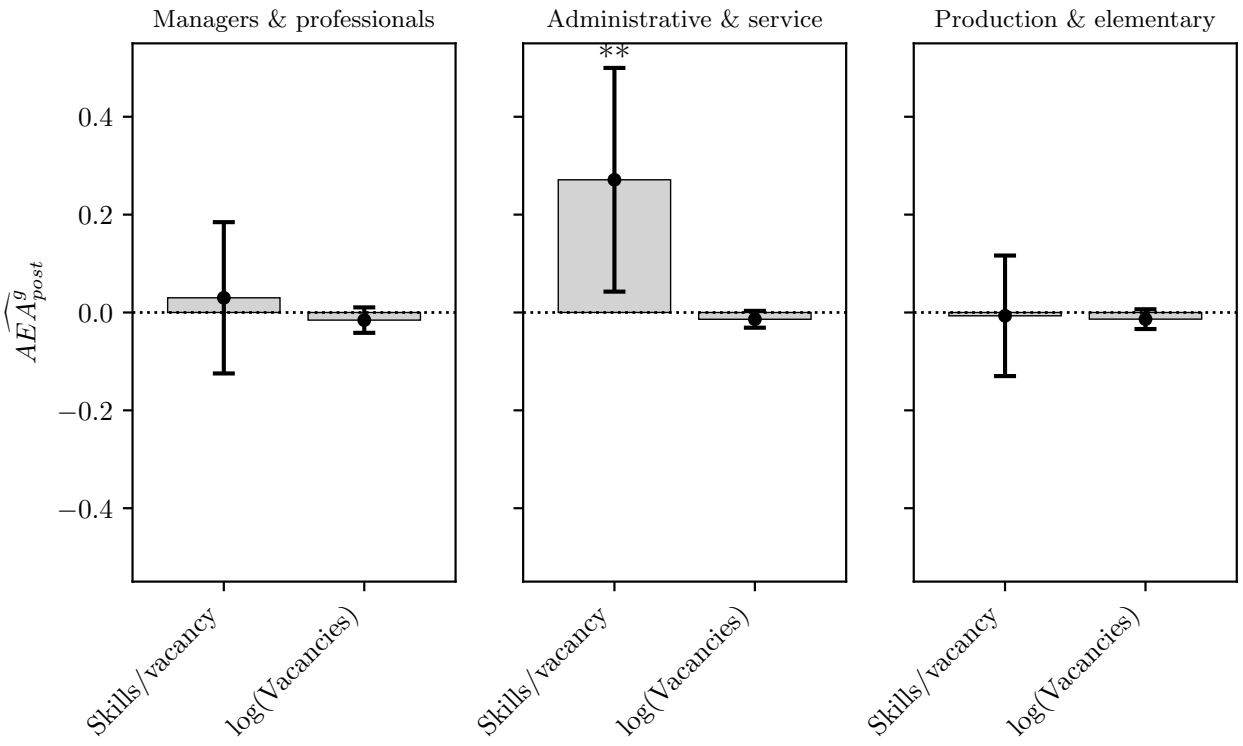
Notes: This figure reports dynamic event-study estimates of the effect of automation on skill demand by occupational group, using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021). Event-time coefficients $\widehat{AEA}_s^g = \sum_c w_c \widehat{AEA}^g(c, c+s)$ are plotted for event time $s \in \{-4, \ldots, 5\}$, where $s = 0$ is the year of the automation event as defined in Equation (7). The estimator uses not-yet-adopter firms as the comparison group. The outcome variables are the shares of soft skills, high-complex hard skills, and low-complex hard skills in total skills demanded within each occupational group. All outcomes are weighted by the number of vacancies. The dotted horizontal lines represent the corresponding aggregate post-treatment estimate $\widehat{AEA}_{post}^g$. Each panel displays estimates for one of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). The lead coefficients ($s < 0$) provide evidence on the parallel trends and no-anticipation assumptions. The sample consists of 3,571 firms; vacancy counts are 33,747 (9,293 obs.) for managers & professionals, 10,497 (3,480 obs.) for administrative & service, and 13,707 (5,899 obs.) for production & elementary. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level.

Figure A5: Effects of Automation on Skill Demand: Controlling for Covariates



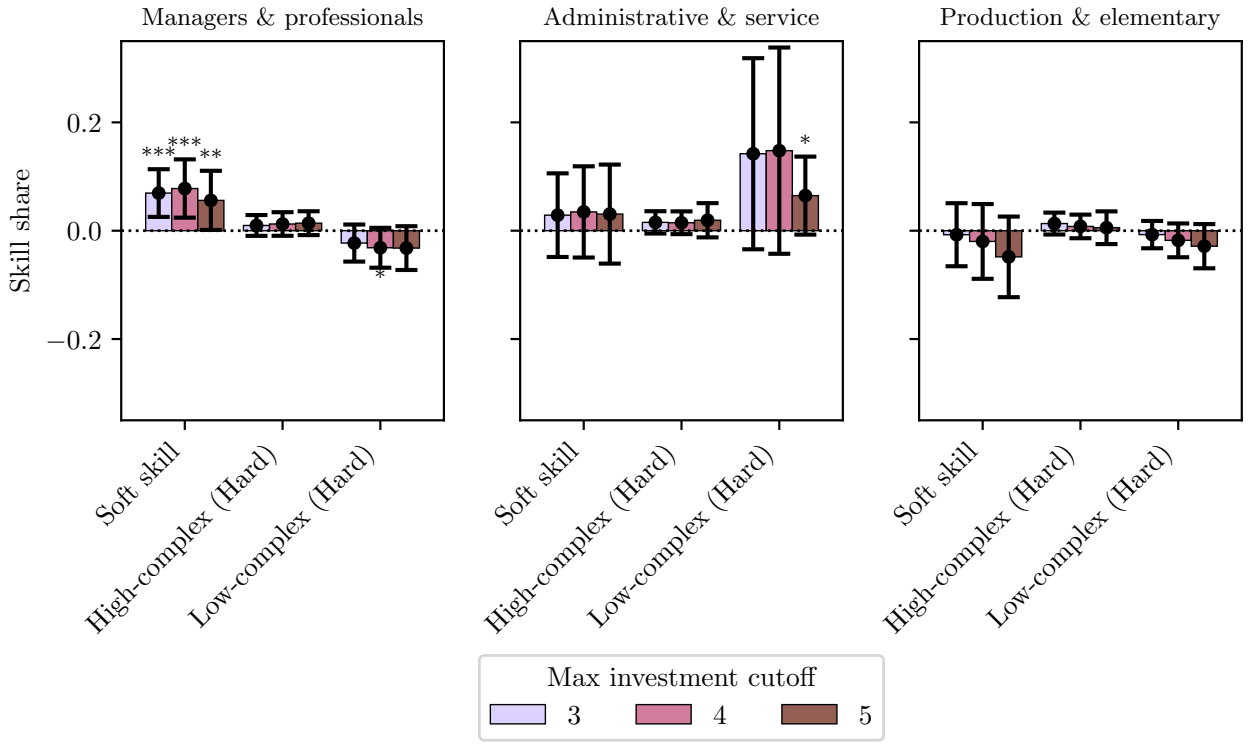
Notes: This figure reports the aggregate post-treatment Average Effect of Automation ($\widehat{AEA}^{g}_{post}$) on skill demand, estimated using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021) as described in Equation (9), with additional pre-automation covariates. The estimator uses not-yet-adopter firms as the comparison group and conditions on: industry (1-digit NACE), quartiles of firm employment, and quartiles of the provincial share of tertiary-educated workers. The outcome variable in each case is the share of a given broad skill type (soft skills, high-complex hard skills, or low-complex hard skills) in total skills demanded within the occupational group. All outcomes are weighted by the number of vacancies. Each panel displays estimates for one of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). The sample consists of 14,051 firm-year observations from 3,571 automating firms, based on 58,728 vacancies. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure A6: Effects of Automation on Vacancies and Skill Intensity



Notes: This figure reports the aggregate post-treatment Average Effect of Automation ($\widehat{AEA}^{g}_{post}$) on vacancy volume and skill intensity, estimated using the staggered difference-in-differences approach of Callaway and Sant'Anna (2021) as described in Equation (9). The estimator uses not-yet-adopter firms as the comparison group. The outcome variables are: the total number of skills identified per vacancy (skills/vacancy), weighted by the number of vacancies; and the natural logarithm of the number of vacancies posted. Each panel displays estimates for one of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). The sample consists of 3,571 firms; vacancy counts are 33,747 (9,293 obs.) for managers & professionals, 10,497 (3,480 obs.) for administrative & service, and 13,707 (5,899 obs.) for production & elementary. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Figure A7: Effects of Automation on Skill Demand: Alternative Event Definitions



Notes: This figure reports the aggregate post-treatment Average Effect of Automation ($\widehat{AEA}^{g}_{post}$) on skill demand under alternative definitions of the automation event threshold in Equation (7). The estimator follows Callaway and Sant'Anna (2021) as described in Equation (9), using not-yet-adopter firms as the comparison group. Cutoff represents different values of automation event threshold $I_{ft} > cutoff * \bar{I}_f$. In the main analysis, a firm's maximum machinery investment must exceed four times its average investment over the panel. Changing the cutoff by one alters the sample of automating firms by approximately 20%. The outcome variable in each case is the share of a given broad skill type (soft skills, high-complex hard skills, or low-complex hard skills) in total skills demanded within the occupational group. All outcomes are weighted by the number of vacancies. Each panel displays estimates for one of the three occupational groups (1-digit ISCO-08): managers & professionals (ISCO 1, 2, 3), administrative & service (ISCO 4, 5), and production & elementary (ISCO 6, 7, 8, 9). The sample size varies with the value of the cutoff, ranging from 2,901 to 4,187 firms. Error bars indicate 95% confidence intervals based on standard errors clustered at the firm level. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

B Proofs of Propositions

This appendix provides proofs for the propositions in Section 2. Throughout, we fix an occupation o and thus, suppress the occupation subscript without loss of generality.

The propositions are derived using a comparative-statics interpretation of the model. Specifically, we compare a pre-automation benchmark (t_0) with a post-automation environment (t_1), without introducing explicit time dependence in the primitives of the model.

At t_0, we consider a benchmark in which automation is not yet adopted or not yet cost-effective for the tasks under consideration. Formally, we assume that all tasks are performed by labor, so that the labor task set is

$$\mathcal{L}_{t_0} = [0, 1].$$

At t_1, changes in technology or relative input costs render a subset of tasks automatable according to the condition

$$\text{Cost}_A(i) < \text{Cost}_L(i),$$

so that the labor task set becomes

$$\mathcal{L}_{t_1} = \mathcal{L} \subseteq [0, 1],$$

with $\mathcal{A} = [0, 1] \setminus \mathcal{L}$ denoting the set of automated tasks.

Thus, the results below compare two environments that differ in the set of tasks performed by labor and, potentially, in the relative weighting of those tasks within occupations.

The propositions characterize distinct comparative-static margins of the model. Proposition 1 isolates the displacement effect arising from the shrinking of the labor task set. Proposition 2 isolates the reweighting effect induced by $\theta_o > 0$ in equation (4). Proposition 3 combines these two channels to characterize relative skill upgrading.

At the pre-automation date t_0, all tasks are assumed to be performed by labor, so the labor task set is

$$\mathcal{L}_{t_0} = [0, 1].$$

At the post-automation date t_1, the labor task set is

$$\mathcal{L}_{t_1} = \mathcal{L} \subseteq [0, 1],$$

where $\mathcal{L} = [0, 1] \setminus \mathcal{A}$ and $\mathcal{A}$ denotes the set of automated tasks.

Pre-automation task weights are given by $\gamma(i) \equiv \gamma_o(i)$, with

$$\int_0^1 \gamma(i)\, di = 1.$$

Post-automation task weights are

$$\gamma'(i) = \frac{\gamma(i) \exp(\theta_o m(i)) \mathbf{1}\{i \in \mathcal{L}\}}{\int_{j \in \mathcal{L}} \gamma(j) \exp(\theta_o m(j))\, dj}.$$

Hence the within-occupation skill share of skill k is

$$S^k_{o,t_0} = \int_0^1 s_k(i)\gamma(i)\, di, \qquad S^k_{o,t_1} = \int_0^1 s_k(i)\gamma'(i)\, di.$$

Equivalently, for any integrable function f,

$$\mathbb{E}_{t_0}[f(i)] \equiv \int_0^1 f(i)\gamma(i)\, di, \qquad \mathbb{E}_{t_1}[f(i)] \equiv \int_0^1 f(i)\gamma'(i)\, di,$$

so that

$$S^k_{o,t_0} = \mathbb{E}_{t_0}[s_k(i)], \qquad S^k_{o,t_1} = \mathbb{E}_{t_1}[s_k(i)].$$

Proof of Proposition 1.

Proposition 1 isolates the displacement channel by abstracting from reweighting across surviving tasks, that is, setting $\theta_o = 0$ in equation (4). In that case,

$$\gamma'(i) = \frac{\gamma(i)\mathbf{1}\{i \in \mathcal{L}\}}{\int_{j\in\mathcal{L}} \gamma(j)\, dj}.$$

Thus post-automation skill demand is the average incidence of skill k among surviving labor tasks:

$$S^k_{o,t_1} = \frac{\int_{i\in\mathcal{L}} s_k(i)\gamma(i)\, di}{\int_{i\in\mathcal{L}} \gamma(i)\, di}.$$

Let

$$\pi_{\mathcal{A}} \equiv \int_{i\in\mathcal{A}} \gamma(i)\, di$$

denote the pre-automation mass of tasks that become automated. Define the average incidence of skill k among automated tasks and surviving tasks, respectively, as

$$\bar{s}^{\mathcal{A}}_k \equiv \frac{\int_{i\in\mathcal{A}} s_k(i)\gamma(i)\, di}{\pi_{\mathcal{A}}}, \qquad \bar{s}^{\mathcal{L}}_k \equiv \frac{\int_{i\in\mathcal{L}} s_k(i)\gamma(i)\, di}{1-\pi_{\mathcal{A}}}.$$

Then

$$S^k_{o,t_0} = \int_0^1 s_k(i)\gamma(i)\, di = \int_{i\in\mathcal{A}} s_k(i)\gamma(i)\, di + \int_{i\in\mathcal{L}} s_k(i)\gamma(i)\, di = \pi_{\mathcal{A}}\bar{s}^{\mathcal{A}}_k + (1-\pi_{\mathcal{A}})\bar{s}^{\mathcal{L}}_k,$$

while

$$S^k_{o,t_1} = \bar{s}^{\mathcal{L}}_k.$$

Therefore,

$$S^k_{o,t_1} - S^k_{o,t_0} = \bar{s}^{\mathcal{L}}_k - \left(\pi_{\mathcal{A}}\bar{s}^{\mathcal{A}}_k + (1-\pi_{\mathcal{A}})\bar{s}^{\mathcal{L}}_k\right) = \pi_{\mathcal{A}}\big(\bar{s}^{\mathcal{L}}_k - \bar{s}^{\mathcal{A}}_k\big).$$

Hence

$$S^k_{o,t_1} < S^k_{o,t_0} \quad \Longleftrightarrow \quad \bar{s}^{\mathcal{A}}_k > \bar{s}^{\mathcal{L}}_k.$$

That is, the within-occupation skill share decreases exactly when the average incidence of

skill k is higher among automated tasks than among surviving tasks. ■

Remark. Proposition 1 characterizes the pure displacement component. Under the full model with $\theta_o > 0$, reweighting across surviving tasks may reinforce or offset this force, but the expression above fully characterizes the sign of the truncation effect.

Proof of Proposition 2.

Fix the surviving task set $\mathcal{L}$ and suppose $\theta_o > 0$. By equation (4),

$$\gamma'(i) \propto \gamma(i) e^{\theta_o m(i)} \mathbf{1}\{i \in \mathcal{L}\}.$$

Thus for any two surviving tasks $i, j \in \mathcal{L}$,

$$\frac{\gamma'(i)}{\gamma'(j)} = \frac{\gamma(i)}{\gamma(j)} e^{\theta_o (m(i) - m(j))}.$$

If $m(i) = 1$ and $m(j) = 0$, then

$$\frac{\gamma'(i)}{\gamma'(j)} = e^{\theta_o} \frac{\gamma(i)}{\gamma(j)} > \frac{\gamma(i)}{\gamma(j)},$$

so complementary tasks receive relatively greater weight post automation than non-complementary surviving tasks. This proves the first claim.

To study the effect on skill demand, partition the surviving task set into

$$\mathcal{L}_1 = \{i \in \mathcal{L} : m(i) = 1\}, \qquad \mathcal{L}_0 = \{i \in \mathcal{L} : m(i) = 0\}.$$

Define

$$A_k = \int_{\mathcal{L}_0} s_k(i) \gamma(i)\, di, \qquad B_k = \int_{\mathcal{L}_1} s_k(i) \gamma(i)\, di,$$

and

$$A = \int_{\mathcal{L}_0} \gamma(i)\, di, \qquad B = \int_{\mathcal{L}_1} \gamma(i)\, di.$$

Then under reweighting parameter $\theta \geq 0$,

$$S^k(\theta) = \frac{A_k + e^\theta B_k}{A + e^\theta B}.$$

Differentiating with respect to θ yields

$$\frac{dS^k(\theta)}{d\theta} = \frac{e^\theta (B_k A - B A_k)}{(A + e^\theta B)^2}.$$

Therefore

$$\frac{dS^k(\theta)}{d\theta} > 0 \quad \Longleftrightarrow \quad \frac{B_k}{B} > \frac{A_k}{A},$$

provided $A, B > 0$. The term B_k/B is the average incidence of skill k among complementary surviving tasks, whereas A_k/A is the corresponding average among non-complementary surviving tasks. Hence, if skill k is more prevalent in complementary tasks than in non-complementary surviving tasks, increasing θ raises the within-occupation skill share of k.

In particular, for $\theta_o > 0$,

$$S^k(\theta_o) > S^k(0)$$

whenever

$$\frac{B_k}{B} > \frac{A_k}{A}.$$

Thus, if complementary tasks rely more heavily on skill k than other surviving tasks, reweighting toward complementary activities increases the within-occupation importance of that skill.

Finally, if no tasks are directly automated, then $\mathcal{L} = [0, 1]$, so the post-automation change operates solely through the tilt

$$\gamma'(i) \propto \gamma(i) e^{\theta_o m(i)}.$$

In that case, any increase in $S^k_{o,t}$ is generated purely by within-occupation reweighting rather than task displacement. ■

Remark. If either $A = 0$ or $B = 0$, then all surviving tasks have the same value of $m(i)$ and reweighting has no effect on relative task weights. In that knife-edge case, $S^k(\theta)$ is constant in θ.

Proof of Proposition 3.

Assume that

$$S^k_{o,t_0} > 0 \qquad \text{and} \qquad S^k_{o,t_1} > 0,$$

so the relative skill demand ratio is well defined. Let

$$\lambda_0 \equiv \frac{S^\ell_{o,t_0}}{S^k_{o,t_0}}, \qquad q(i) \equiv s_\ell(i) - \lambda_0 s_k(i).$$

By construction,

$$\int_0^1 q(i)\gamma(i)\, di = S^\ell_{o,t_0} - \lambda_0 S^k_{o,t_0} = 0.$$

Moreover,

$$\int_0^1 q(i)\gamma'(i)\, di = S^\ell_{o,t_1} - \lambda_0 S^k_{o,t_1}.$$

Hence

$$\frac{S^\ell_{o,t_1}}{S^k_{o,t_1}} \geq \frac{S^\ell_{o,t_0}}{S^k_{o,t_0}} \quad \Longleftrightarrow \quad \int_0^1 q(i)\gamma'(i)\, di \geq 0.$$

To obtain this inequality, impose the following monotone-shift condition: automation shifts relative task weight toward tasks with weakly higher values of $q(i)$, in the sense that the post-automation density γ' first-order stochastically tilts weight toward tasks for which skill ℓ is relatively more important than skill k. A sufficient condition is that $q(i)$ is weakly higher in tasks with lower automatability and in complementary tasks ($m(i) = 1$), so that both channels of automation act in the same direction:

1. tasks with relatively high $q(i)$ are less likely to be removed from the labor task set; and
2. among surviving tasks, those with relatively high $q(i)$ receive weakly greater weight when $\theta_o > 0$.

Under this condition, the post-automation task distribution γ' places weakly more weight on tasks with higher values of $q(i)$ than the pre-automation distribution γ. Therefore,

$$\int_0^1 q(i)\gamma'(i)\,di \geq \int_0^1 q(i)\gamma(i)\,di = 0.$$

It follows that

$$S^{\ell}_{o,t_1} - \lambda_0 S^{k}_{o,t_1} \geq 0,$$

which is equivalent to

$$\frac{S^{\ell}_{o,t_1}}{S^{k}_{o,t_1}} \geq \frac{S^{\ell}_{o,t_0}}{S^{k}_{o,t_0}}.$$

Thus, under monotone reweighting toward tasks that are relatively more intensive in skill ℓ than in skill k, automation weakly increases the relative within-occupation demand for ℓ.

■

Remark. Proposition 3 combines the logic of Propositions 1 and 2. Relative skill upgrading arises when automation disproportionately removes or downweights tasks intensive in skill k and preserves or upweights tasks intensive in skill ℓ.

Derivation of Equation (6)

Starting from equation (5),

$$S^{k}_{t_1} - S^{k}_{t_0} = \sum_o w_{o,t_1} S^{k}_{o,t_1} - \sum_o w_{o,t_0} S^{k}_{o,t_0}.$$

Using

$$w_{o,t_1} = w_{o,t_0} + \Delta w_o, \qquad S^{k}_{o,t_1} = S^{k}_{o,t_0} + \Delta S^{k}_{o},$$

we obtain

$$\Delta S^{k} = \sum_o (w_{o,t_0} + \Delta w_o)(S^{k}_{o,t_0} + \Delta S^{k}_{o}) - \sum_o w_{o,t_0} S^{k}_{o,t_0}.$$

Expanding the product yields

$$\Delta S^k = \sum_o \Delta w_o \, S^k_{o,t_0} + \sum_o w_{o,t_0} \, \Delta S^k_o + \sum_o \Delta w_o \, \Delta S^k_o.$$

Rearranging terms gives equation (6):

$$\Delta S^k = \underbrace{\sum_o \Delta w_o \, S^k_{o,t_0}}_{\text{Between-Occupation Change}} + \underbrace{\sum_o w_{o,t_0} \, \Delta S^k_o}_{\text{Within-Occupation Change}} + \underbrace{\sum_o \Delta S^k_o \Delta w_o}_{\text{Interaction}}.$$

This is an accounting identity.

C Data sample

The following sample restrictions are imposed on the firms. The main period of our analysis is 2010-2021. This is limited in the lower end by the update to the occupational classification from DISCO-88 to DISCO-08, which makes a within-occupation analysis including this changing period more difficult (The job vacancy data is available from 2008, meaning two years are dropped from the analysis). The upper bound of the sample period is limited to 2021, this being the last year of available register data.

In the job vacancy data, some job titles lack a formal job title. These vacancies are removed, as they cannot be used in the within-occupation analysis.

We impose the following sample limitations on the main panel of firms: As our data on automation consists of machinery investments, we drop firms in sectors with no reported machinery investments at all, along with financial sectors, and sectors where machinery is more likely to refer exclusively to infrastructure such as pumps and heating. We impose a size restriction of firms having at least three employees in the entire time period, so as to not confuse the effect of automation with a new firm building up their capital. Further, we require that a firm has to be active in at least three years. Next, some small firms are not covered by the balance sheet data used to identify automation events, and so they are dropped. After this, we limit the sample to firms active after 2010, for the reason mentioned in the start of this section. Next, we exclude any firms that did not post a single vacancy in the entire sample period, as these would not provide any value for the main event studies. Finally, as our event studies uses not-yet-treated units as the control group, we remove firms with no automation events. Table A8 displays how each stage of sample restrictions affect the sample size.

TABLE A8
Sample Construction

Description	Remaining Firms
Original sample	1,555,818
Remove irrelevant sectors	1,004,946
Limit to firms with ≥ 3 employees in full active period	40,004
Keep firms with ≥ 3 active years	34,685
Drop firms without investment data	31,807
Limit sample to post-2010	30,104
Drop zero vacancy firms	16,289
Automating firms	4,982

Notes: This table documents the sequential sample restrictions applied to construct the analysis sample, reporting the number of firms remaining after each step. The original universe comprises all firms in the Danish register data. Restrictions include: removal of sectors with no recorded machinery investment, financial sectors, or sectors where machinery likely reflects infrastructure rather than automation; a minimum firm-size requirement; a minimum activity requirement; exclusion of firms without balance-sheet investment data; restriction to firms active after 2010, when the DISCO-08 occupational classification came into effect; removal of firms that never posted a vacancy over the sample period; and, finally, identifying firms experiencing an automation event as defined in Equation (7). The main analysis sample consists of 4,982 automating firms.