I Z A Institute of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 16740

De-Routinization in the Fourth Industrial Revolution - Firm-Level Evidence

Melanie Arntz
Sabrina Genz
Terry Gregory
Florian Lehmer
Ulrich Zierahn-Weilage

JANUARY 2024

I Z A Institute
of Labor Economics
Initiated by Deutsche Post Foundation

DISCUSSION PAPER SERIES

IZA DP No. 16740

De-Routinization in the Fourth Industrial Revolution - Firm-Level Evidence

Melanie Arntz
ZEW and University of Heidelberg

Sabrina Genz
Utrecht University, IAB and TPRI

Terry Gregory
LISER, ZEW and IZA

Florian Lehmer
IAB

Ulrich Zierahn-Weilage
Utrecht University and ZEW

JANUARY 2024

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

De-Routinization in the Fourth Industrial Revolution - Firm-Level Evidence*

This paper examines the extent to which aggregate-level de-routinization can be attributed to firm-level technology adoption during the most recent technological expansion. We use administrative data and a novel firm survey to distinguish frontier technologies from older technologies. We find that adopters of frontier technologies contribute substantially to deroutinization. However, this is driven only by a subset of these firms: large adopters replace routine jobs and less routine-intensive adopters experience faster growth. These scale and composition effects reflect firms' readiness to adopt and implement frontier technologies. Our results suggest that an acceleration of technology adoption would be associated with faster de-routinization and an increase in between-firm heterogeneity.

JEL Classification: J21, J23, J24, O33

Keywords: technology, automation, tasks, capital-labor substitution, decomposition

Corresponding author:
Terry Gregory
Luxembourg Institute of Socio-Economic Research: LISER
11 Prte des Sciences
4366 Esch-Belval Esch-sur-Alzette
Luxemburg
E-mail: terry.gregory@liser.lu

* We thank Steven Dhondt, Christina Gathmann, Michael Handel, Attila Lindner, Piotr Lewandowski, Ludivine Martin, Emilie Rademakers, Pascual Restrepo, Anna Salomons, Ralf Wilke, and Nicolas Ziebarth for helpful suggestions. Furthermore, we thank participants of the TASK VI Conference on "the Digital and Ecological Transformation of the Labour Market", the EALE 2023 conference, the UNTANGLED workshop on "Skill Changes and the Value of Skills", and seminar participants at Boston University, Utrecht University, Heidelberg University, Luxembourg University, IAB, LISER, and ZEW for comments. We finally thank the German Federal Ministry of Education and Research (BMBF) as well as the Instituut Gak for financial support. We also gratefully acknowledge financial support from the Leibniz Association through the Leibniz Professorship for Applied Labor Economics at the University of Heidelberg (P56/2017). This project has received further funding from the European Union's Horizon 2020 research and innovation program under grant agreement number 101004494. The contents of this publication are the sole responsibility of the authors and do not necessarily reflect the opinion of the European Union.

1 Introduction

An extensive body of literature shows that computerization, beginning in the 1970s, has replaced workers in routine codifiable tasks (Autor et al., 2003; Acemoglu and Autor, 2011). This phenomenon, known as Routine Replacing Technological Change (RRTC), has led to a polarization of employment by reducing the share of middle-wage routine jobs while simultaneously raising the shares of both high-paid non-routine cognitive and low-paid service jobs (Goos et al., 2009, 2014; Cortes, 2016). Recently, the rapid evolution of new technologies, particularly artificial intelligence (AI) and other technologies has sparked questions about whether these emerging frontier technologies might substitute increasingly non-routine tasks (Brynjolfsson et al., 2018; Webb, 2020). Very few studies are able to document the firm-level adoption and diffusion of the latest frontier technologies such as AI (McElheran et al., 2022; Zolas et al., 2020; Genz et al., 2021). Moreover, the implications of technology adoption for the aggregate occupational composition remain unclear as existing studies on the effects of technology adoption tend to focus on structural shifts within firms (Gaggl and Wright, 2017; Aghion et al., 2020). Yet, whether aggregate-level de-routinization stems from the displacement of routine jobs within firms adopting frontier technologies, or from differential trends in employment growth or de-routinization among heterogeneous firms, comes with different implications for the routine-replacing character of frontier technologies and related between-firm heterogeneity (Barth et al., 2016; Song et al., 2019; Cortes et al., 2023). As Seamans and Raj (2019) have pointed out, a significant challenge to investigating these issues is the lack of sufficient firm-level data with information on the adoption of frontier technology.

To fill this knowledge gap, we collect novel data on the actual adoption of frontier technologies in the German economy (as opposed to technology exposure[1]) and examine how the diffusion of these technologies contributes to aggregate occupational changes. For this, we decompose aggregate changes in occupational employment into the contribution by firms adopting frontier technologies compared to those firms adopting more mature technologies. We further examine whether these occupation changes take place within or between firms, hence shedding light on the role of heterogeneous firms. Our paper is thus the first to provide a full picture of how the diffusion of frontier technologies in the entire economy affects aggregate occupational changes.

[1]A series of papers use technology exposure measures to assess the impact of frontier technologies on the labor market (Acemoglu et al., 2022b; Brynjolfsson et al., 2018; Gathmann and Grimm, 2023; Webb, 2020). In contrast to those, we focus on actual technology adoption.

Our database consists of novel survey data on firm-level adoption of frontier technologies among 2,032 German manufacturing and service firms, linked to administrative social security records of all workers employed at the surveyed firms. The data allow us to distinguish between non-investing firms, firms that mainly invest in digital technologies, such as computers and robots, and firms that mainly invest in frontier technologies such as AI, augmented reality, smart factories, or cloud computing. With the administrative worker records including detailed occupational information, we are able to track the firms' workforce composition across time and link it to the firm's technology adoption status.

We first document that both digital and frontier adopters have a larger share of skilled workers, and are more likely to be service firms than non-adopters. In line with the RRTC literature, the initial occupation structure generally matters for modern technology adoption. Surprisingly, these factors do not explain whether firms invest in frontier or digital technologies. Instead, high-speed internet access is a key predictor of adopting frontier technologies.

We then study the aggregate relationship between technology adoption and changing employment structures. We find a strong de-routinization of the German workforce between 2011 and 2016: The share of routine jobs declined, while non-routine cognitive jobs rose strongly, and also non-routine manual jobs increased slightly. Importantly, the share of routine occupations fell the most in the group of firms adopting frontier technologies, followed by those adopting digital technologies. To unravel whether the aggregate changes arise within or between firms, we apply the decomposition approach by Autor et al. (2020) and Acemoglu et al. (2020) to shifts in employment shares and distinguish three groups: non-adopters, digital adopters, and frontier adopters.[2] The results reveal a surprising pattern: The substitution of routine workers by technology does not take place among all frontier adopters. In fact, the average frontier firm (insignificantly) increases its share of routine workers. Likewise, de-routinization cannot be explained by initial differences between technology adopters and non-adopters. Instead, de-routinization is driven by a subset of frontier firms: Initially larger frontier adopters experience faster de-routinization (scale effect), and frontier adopters with an initially less routine-intensive workforce grow faster (composition effect).

Finally, we delve deeper into potential explanations for these observed heterogeneities using firm-level estimations to study what explains these scale and composition effects.

[2] Autor et al. (2020) and Acemoglu et al. (2020) focus on shifts in labor shares. We augment their decomposition by introducing sampling weights to account for the firm survey design, we decompose the residual changes in the covariance term, and we distinguish between groups rather than two.

Interestingly, neither firm characteristics, worker characteristics, nor firms' capital structure explain the scale and composition effect. Instead, the scale effect appears to be related to how firms perceive changing skill demands and related training needs. This suggests that larger firms are better prepared to recognize and make complementary investments in upskilling workers when adopting frontier technologies. The composition effect is extremely robust to adding different sets of covariates, suggesting that firms' initial occupation structure plays a crucial role irrespective of other observable characteristics. Adopters initially specialized in non-routine cognitive jobs grow faster, indicating that firms' occupational specialization matters for the firm's ability to benefit from new frontier technologies. Our findings support the notion that the successful adoption of frontier technologies requires complementary investments into workers' skills and restructuring (Brynjolfsson and Hitt, 2003; Harrigan et al., 2021).

Our paper conveys two major takeaways: First, the diffusion of frontier technologies is likely to foster aggregate de-routinization, as de-routinization is faster among frontier technology adopters. Second, this de-routinization is likely to be driven by a subset of frontier firms that are best prepared to benefit from these technologies. This might also indicate that the speed of de-routinization would be even faster if more firms had the means to make complementary investments in workers' skills.

Our paper is related to three literature strands. Firstly, we contribute to the long-standing literature on the response of workforce structures to *aggregate*-level technology shocks, measured at the industry, region, or country level. In particular, the literature on RRTC shows that computerization (which is comparable to the adoption of our digital technologies) induces a decline in the overall share of workers in routine occupations in the economy (see, e.g., Autor et al., 2003; Goos et al., 2009; Acemoglu and Autor, 2011; Autor and Dorn, 2013; Goos et al., 2014). A related literature focuses on the adoption of industrial robots (which is a specific digital technology in the manufacturing sector) and documents employment shifts across industries and regions (see, e.g., Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Dauth et al., 2021). Compared to this literature, we are the first to provide micro-founded evidence on how the diffusion of frontier technologies in all parts of the economy affects the aggregate occupational composition. Our findings reveal that de-routinization is far from a uniform process in firms adopting frontier technologies, but hinges very much on conducive conditions that need to be in place when adopting such technologies.

Secondly, we contribute to a recent literature focusing on workforce structures in response to *firm*-level technology adoption. Corresponding studies examined the impact of firm-level ICT investments (Bartel et al., 2007; Böckerman et al., 2019; Gaggl and Wright, 2017; DeStefano et al., 2018), robot adoption (Acemoglu et al., 2020; Bonfiglioli et al., 2020; Dixon et al., 2021; Humlum, 2019; Koch et al., 2021) and investments into automation technologies more generally (Aghion et al., 2020; Bessen et al., 2020). Recent studies look at the impact of artificial intelligence (Acemoglu et al., 2022a,b; Zolas et al., 2020; McElheran et al., 2022). Overall, these studies find that firm-level technology adoption enforces organizational change that coincides with labor demand shifts towards non-routine, more skilled, and better-paid workers, but also higher inequality between workers within the firm.[3] We contribute to this literature by differentiating between older digital technologies and frontier technologies and by focusing on firms from both manufacturing and service industries. We thereby show that the routine-replacing effect of frontier technologies hinges on firm-specific rather than sector-specific factors.

Finally, we contribute to the emerging literature on firm-level evidence on sources of inequality. Studies show that employment reallocation within and between firms contributes to the changing occupational structure, referred to as job polarization (Kerr et al., 2020; Harrigan et al., 2021). Other research argues that technological change leads to higher employment concentration in more productive firms and an increase in between-firm wage inequality (Cortes et al., 2023). Further studies show that the falling labor share is driven by employment reallocation between firms, referred to as the superstar phenomenon (Autor et al., 2020; Acemoglu et al., 2020). We contribute to this literature by documenting employment reallocation relating to frontier technology adoption. We can do so because our data allows us to differentiate employment reallocation between firms that take place within versus between adoption groups. Thereby we show that job polarization is driven by a subset of frontier technology adopters, giving rise to a new dimension of workplace polarization within the group of adopters.

The remainder of this paper proceeds as follows. Section 2 describes the data and develops our technology adoption measure. Section 3 provides descriptive statistics on adopters and non-adopters and studies firm-level determinants of technology adoption. Section 4 relies on a decomposition analysis to document a link between aggregate de-routinization and firm-level technology adoption, showing that heterogeneity between adopters of frontier technologies plays a crucial role. Section 5 delves into mechanisms behind this heterogeneity and section 6 concludes.

[3] While we concentrate on workforce structures, further studies show positive impacts on firm performance (Cirillo et al., 2022; Bettiol et al., 2019), productivity (Cathles et al., 2020) and innovation activity (Babina et al., 2022; Rammer et al., 2022).

2 Data

2.1 Data sources

For our analysis, we link a representative firm survey on technology adoption with employment biographies from social security records for all workers employed in the interviewed firms.

Firm survey on technology adoption. We conducted a survey among a representative sample of 2,032 German establishments between March and May 2016, the "IAB-ZEW-Labor Market 4.0-Establishment Survey (BIZA)" to identify the adoption of technologies of the fourth industrial revolution.[4] The establishments in our survey are a stratified random sample of all German establishments with at least one employee subject to social security contributions which are registered at the German Federal Employment Agency. The German administrative data contains information on single-location establishments and does not contain information on whether these establishments belong to a multi-site company or not. For simplicity, we refer to these establishments as firms in the remainder of the paper.

The interviews were held with either the production managers or the firms' general manager and lasted an average of 30 minutes. The survey covers three main topics: (1) the relevance of and perceptions about frontier technologies; (2) the sophistication level of a firm's technology; and (3) the firms' demand for skills and competencies. We collect all information contemporaneously for 2016 and ask for retrospective information for 2011 for several key questions. In particular, we use retrospective information to track the adoption of frontier technologies.

All firms received a written invitation for the survey before firms were called to conduct a CATI-interview. Among all contacted firms, 2,032 firms completed the survey, yielding a response rate of 31.5 percent. Non-participation was typically due to general refusal of interviews, especially over the phone, or due to time constraints. The use of modern technologies played only a minor role in non-participation decisions. Genz et al. (2021) provide a detailed analysis of non-participation for the dataset, showing that the surveyed firms are comparable in terms of size, industry, and location to the population of German firms (see Appendix A.1 in Genz et al. (2021)). Hence, our survey is representative of German firms, and potential biases concerning technology adoption, if any, are likely to be small.

[4] The data can be accessed at the Research Data Centre (FDZ) of the German Federal Employment Agency, see `https://fdz.iab.de/en/our-data-products/establishment-data/biza/`

Weights. Our sample is stratified by firm size (four categories), region (East/West Germany), and industrial sector (five categories)[5] and covers both service and manufacturing firms. To ensure sufficient observations, we conducted at least 50 interviews within the resulting 40 cells.[6] This naturally leads to oversampling of certain cells (relative to the overall population of firms). We correct for oversampling by computing firm stratification weights w_f as the inverse inclusion probability of firms in our survey. Weights are scaled to a mean value of one, such that the sum of weights reflects the number of firms.[7] We apply the firm stratification weights w_f in empirical analyses which are representative of all German firms, for instance, when we focus on the changes of average employment shares across task domains within firms.

For analyses with a focus on workforce composition, or when firm size matters, we rely on employment weights $\tilde{s}_f$. In particular, we use the firm stratification weights w_f to compute the employment weights $\tilde{s}_f = \frac{w_f s_f}{\sum_{f' \in F} w_{f'} s_{f'}}$, where s_f is a firm's share of total employment. We scale these weights to a mean of one analogous to the firm stratification weights. We apply the employment-weighted firm stratification weights $\tilde{s}_f$ in empirical analyses which are representative of the German workforce, for instance, when we analyze changes in the aggregate employment shares across task domains.[8]

Employment histories. We link our survey data to employment biographies from social security records (IAB Beschäftigtenhistorik (BeH) V10.03.00, Nürnberg 2018) for all workers employed in the surveyed firms between 2011 and 2016. The social security system covers roughly 80% of the German workforce, omitting information about the self-employed and civil servants. The social security records include, among other items, information on workers' employment status, occupation, industry, and earnings.[9] We transform the data for all workers between 18 and 65 years of age in regular jobs subject to social security contributions to an annual panel between 2011 and 2016.[10] This implies that the data excludes minor employment, apprentices, and family workers, which leaves 1,191,734 observations for 318,316 employees between 2011 and

[5]See Data Appendix B.1 for details.

[6]We merge the cells for firms with "50-200 employees" and "200 and more employees" in the the East German ICT sector due to the small number of large ICT firms in East Germany.

[7]The sample of firms was drawn in 2014, weights therefore are representative of the 2014 distribution of firms and are time constant.

[8]Depending on the type of analysis, we use either time-varying employment weights based on time-varying employment information or time-constant employment weights based on firms' initial employment in 2011.

[9]Wages are reported only up to the social security contribution limit. We impute wages using Tobit regressions following Card et al. (2013) and Dustmann et al. (2009).

[10]We follow the standard procedure in the literature using German administrative data and rely on employment spells that overlap June 30th of each year.

2016.

Based on this sample of workers, we calculate the firm's occupation structure and other firm-level indicators such as total employment and average wages. For the occupational composition, we use the main task categorization for occupations by Dengler et al. (2014). They exploit information on competencies and skills from the German expert database BERUFENET of the Federal Employment Agency to calculate the main task type of occupations. We merge the main task type to our data via the occupation code (KldB-2010, 3-digit) to classify occupations and calculate the firms' annual workforce size by task domain, thereby distinguishing between occupations that predominantly use (1) non-routine cognitive (2) routine cognitive and manual, and (3) non-routine manual tasks. This enables us to detect employment shifts from routine occupations to non-routine occupations over time.[11] The occupational classification (KldB-2010, 5-digit) furthermore contains information on job requirement categories that are sorted by increasing complexity. It distinguishes (1) unskilled workers, (2) professionals, (3) specialists, and (4) experts. We prefer this skill measure to a formal education measure because it captures the skills required to perform a job from the employer's perspective. We also calculate the number of workers by age as of 2011 and distinguish between (1) younger than 40, (2) between 40 and 55, and (3) older than 55 years of age.

High-speed internet access. We use data on firms' access to high-speed internet in the determinant analysis of technology adoption in section 3. For this, we use the kilometer distance of the firm to the next main distribution frame (MDF). In Germany, the so-called "last mile" between the firm location and the MDF typically relies on copper wires, which leads to reductions in internet speed with distance, as highlighted by Falck et al. (2014, 2021). While these studies used the distances between the centroids of all municipalities to the next MDF, we measure the exact distance between the establishment location and the closest MDF.

2.2 Sample selection

Since we aim to examine occupational employment shifts across three task domains at the firm level between 2011 and 2016, we impose several restrictions on the sample of 2,032 surveyed

[11] Our measure classifies occupations according to their main task in 2011 and does not capture variation in task content within occupations across time, which has been highlighted as an additional dimension of occupational change (Acemoglu and Autor, 2011; Spitz-Oener, 2006; Freeman et al., 2020).

establishments.[12] First, we restrict our sample to firms that exist in all years, i.e., that were founded prior to June 30th, 2011, and that are observable in our administrative data until 2016 (-235 firms). In addition, we impose the restriction that these firms have at least one employee subject to social security contributions in regular employment on June 30th, 2011 and 2016 (-122 firms). A further 15 firms are dropped due to missing information on the technology indicators which are at the core of our analysis. These restrictions leave us with 1,660 firms that are observable in our administrative data for the entire observation period, have at least one employee across time, and for whom we have information on technology adoption between 2011 and 2016.

2.3 Defining technology adoption

The key advantage of our survey is that it contains a direct firm-level measure of technology adoption. In particular, our survey measures the firms' level of technological sophistication in 2016 and 2011 and uses this survey information to classify firms into three adopter groups.

Surveying a firm's level of technological sophistication. In order to survey the firms' level of technological sophistication, we introduce respondents to the concept of three technology levels that differ in their degree of automation and digitalization and provide an example for each of these levels, see Table 1. In particular, we distinguish frontier technologies of the fourth industrial revolution, computer-controlled technologies of the third industrial revolution, and older technologies belonging to the first and second industrial revolutions. The definitions of each of these technology levels are kept rather generic to ensure that the distinction applies to and is comparable across all firms in our sample. To facilitate a common understanding, we give examples of technologies and provide respondents with further examples upon request. Moreover, we measure the level of technology separately for electronic office and communication (O&C) and production (PROD) tools to allow for more targeted questions and examples throughout the survey. This separate assessment of the firm's technology level also accounts for the fact that technological progress need not be implemented at the same pace in both domains.[13]

Manual work tools refers to the lowest technology level and comprises old technology associated

[12] There is a structural break in the occupational classification in 2011 that has been harmonized but still might limit the comparability of the task structure across time. Appendix E shows that the results are robust to using 2012 rather than 2011 as the base year for our analysis.

[13] Since our firm survey includes firms in the service sector, not all surveyed firms use production tools. Firms without production tools (61 percent) assessed only the level of technology of their office and communication tools.

Table 1: Firms' level of technological sophistication

Technology level $g \in \{1.0/2.0, 3.0, 4.0\}$	Production tools (k_g^{PROD})	Office and communication tools ($k_g^{O\&C}$)	Work process organization
Frontier technology (fourth industrial revolution, $g \in \{4.0\}$)	**Self-controlled** ($k_{4.0}^{PROD}$) e.g., production facilities up to smart factories, cyber-physical systems and internet of things	**IT-integrated** ($k_{4.0}^{O\&C}$) e.g., analytic tools for big data, cloud computing systems, internet platforms such as Amazon, shop systems or online-markets	Technology performs work process largely automatically and autonomously
Digital technology (third industrial revolution, $g \in \{3.0\}$)	**Indirectly controlled** ($k_{3.0}^{PROD}$) e.g., CNC machines, industrial robots or process engineering systems	**IT-supported** ($k_{3.0}^{O\&C}$) e.g., computers, terminals, electronic checkout systems or CAD-systems	Humans are indirectly involved in work process
Manual technology (first and second industrial revolution, $g \in \{1.0/2.0\}$)	**Manually controlled** ($k_{1.0/2.0}^{PROD}$) e.g., drilling machine, motor vehicles or X-ray machine	**Not IT-supported** ($k_{1.0/2.0}^{O\&C}$) e.g., telephones, fax and copy machines	Humans largely conduct work process manually

Notes: The table shows the technology levels as they were defined and explained (including examples) to the interviewees during the CATI-interview. The interviewees were then asked to give an estimate of how their tools are divided among these technology classes (in percent). Technology levels refer to the four industrial revolutions.

with the first and second industrial revolution which requires humans to be largely involved in the work processes. Production tools in this technology class is manually controlled, such as drilling machines, motor vehicles, or X-ray machines. Office and communication tools in this class is not IT-supported, such as analog telephones, faxes, or copy machines.

The intermediate technology level encompasses digital technologies associated with the third industrial revolution, which have been in place since the 1970s. It refers to tools where humans involvement in work processes is indirect. These technologies are supported by IT to automate specific sub-processes and facilitate new applications. Production tools in this class is indirectly controlled, such as CNC machines, industrial robots, or process engineering systems. Similarly, office and communication tools in this class are IT-supported, including computers, terminals, electronic checkout systems, or CAD systems.

Technologies at the highest level, representing frontier technologies associated with the fourth industrial revolution, are fully integrated into the firms' central IT systems, facilitating largely automated work processes without human intervention. This class includes self-controlled production tools, such as smart factories, cyber-physical systems, and the Internet of Things. Moreover, it encompasses IT-integrated office and communication tools, such as the use of big data, cloud computing, and online markets. The highest technology level hence covers the most recent technologies that we consider to reflect the technological frontier as of 2016.

After providing interviewees with an overview of these technology classes, we ask them to separately assess the share (in percent) of the firm's production and O&C tools that is currently based on each of these in 2016. We also ask for a corresponding retrospective assessment for the technology classes in 2011.[14]

Technology-specific investments. Based on these shares, we calculate the overall share of firms' tools by technology level g, where $g \in \{1.0/2.0, 3.0, 4.0\}$, at time t in firm f as follows:

$$k_{fgt} = \Big((1 - \delta_f)\; k_{fgt}^{O\&C} + \delta_f\; k_{fgt}^{PROD}\Big) \tag{1}$$

where δ_f equals the time-constant initial-level share of blue-collar workers in the firm that we derive from the administrative employment histories that can be linked with our survey.[15] For simplicity, we refer to firms' tools by technology level as technology shares k_{fgt}. Technology shares add up to unity separately for production and office and communication tools within each firm and year.

While changes in technology shares are informative about the composition of firms' technology, they contain limited information on firms' technology investments since changes in technology shares do not capture those technology investments that yield the relative composition unaltered. To capture differences in the total investment sums, we combine the information on the technology shares k_{fgt} with a measure of the firm's capital *stock* K_{ft}, and derive a technology-specific capital stock across time as $K_{fgt} = k_{fgt} \times K_{ft}$. Since we do not observe information on the absolute value of the capital stock at the firm level, we impute capital stock information K_{ft} for the firms in our sample from the IAB Establishment Panel, which requires additional assumptions.[16] The obtained technology-specific capital stocks K_{fgt} allow us to calculate the firm's net investments I_{fg} into technology-specific capital stocks by taking the five-year difference between 2011 and 2016:

$$I_{fg} = \Delta K_{fg} = K_{fgt=2016} - K_{fgt=2011} \tag{2}$$

[14] Column 1 in Appendix Table A.1 reports the average shares across all responses obtained from the firm survey. Note that Genz et al. (2021) show plausible patterns when comparing the technology shares from the survey against existing measures at the aggregate industry level, such as ICT capital shares from EU KLEMS data.

[15] The classification is based on 1-digit levels of International Standard Classification of Occupations (ISCO 2008): white-collar (ISCO codes 1 to 5); blue-collar (ISCO codes 6 to 9).

[16] See Appendix B.2 for details. Given that the survey was conducted as a CATI, we considered it infeasible to ask for the technology-specific capital stocks across time.

Classification of adopter groups. We use these net investments I_{fg} to classify firms into three adopter groups depending on their main investments from 2011 to 2016:

$$g\text{-adopter} \iff I_{fg} \quad = \quad max\left(I_{f,g=1.0/2.0(non)}, I_{f,g=3.0(digital)}, I_{f,g=4.0(frontier)}\right) \qquad (3)$$

Hence, *non-adopters* mainly invest in older manual technologies, *digital adopters* mainly adopt computer-controlled digital technologies and *frontier adopters* mainly invest in autonomous cutting-edge technologies. A firm belongs to the adopter group $g \in (1.0/2.0(non); 3.0(digital); 4.0(frontier))$ (1) if the firm raises its share of g-technologies more than the other shares, irrespective of whether the firm expands or shrinks its capital stock, (2) if the firm leaves the relative technology shares unchanged, but expands its overall capital stock and used primarily g-technologies in 2011, (3) if the firm raises its overall capital stock and simultaneously raises its g-technology share. Hence, our definition considers both initial stocks and changes in technology shares when classifying firms' adoption status. Appendix B.3 provides a formal definition and discusses special cases.

Note that we prefer to classify firms based on the imputed technology-specific changes in capital stocks ΔK_{fg} rather than using the reported changes in technology shares Δk_{fg}. Assigning firms to adoption groups based on technology-specific changes in capital stocks ΔK_{fg} emphasizes firms' initial technology composition, yielding a more stringent classification of frontier adopters.[17]

3 Descriptive statistics on technology adoption

Before analyzing the role of adoption groups for the aggregate occupational composition in Germany, we first show the technology adoption trends between 2011 and 2016. Second, we examine the selection of firms into the three adoption groups to understand what predicts firms' investments into frontier technologies and to detect heterogeneity in the initial workforce composition between adoption groups.

3.1 Technology use and adoption in Germany

Table 2 shows that our final sample consists of 525 non-adopters, 962 digital adopters, and 173 frontier adopters when applying the aforementioned classification. Frontier adopters only

[17] In Appendix B.3 we discuss the strengths and weaknesses of both approaches and provide illustrative examples of the assignment into adoption groups. In Appendix E, we show a replication of our main results using the alternative approach to assign adoption groups by focusing on changes in technology shares.

represent 7.9% of all German firms, but 10.4% of all workers in 2011 (see Panel C). Thus, only a small fraction of frontier firms predominantly invest in the most recent technologies such as augmented reality, smart factories, or cloud computing, while most German firms focus on either computer-controlled digital technologies or even older manual technologies.

Table 2: Technology use and employment growth across German firms

	All		non-adopters		digital adopters		frontier adopters	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Technology-specific capital stock in 2011 (in 1,000 €) and in %								
manual tech	72.6	46.9%	59.0	71.2%	82.3	40.2%	88.9	35.2%
digital tech	71.7	46.3%	20.1	24.3%	112.5	54.9%	109.9	43.5%
frontier tech	10.7	6.9%	3.7	4.5%	10.0	4.9%	53.9	21.3%
Panel B: Technology-specific capital stock in 2016 in 1,000 €) and in %								
manual tech	111.7	38.4%	116.1	70.9%	110.9	28.7%	91.4	21.7%
digital tech	148.7	51.1%	40.4	24.7%	246.0	63.7%	160.8	38.2%
frontier tech	30.6	10.5%	7.1	4.4%	29.4	7.6%	168.5	40.0%
Panel C: Representing number of firms and workers								
# firms (in K)	1925.1	100.0%	848.4	44.1%	924.7	48.0%	151.9	7.9%
# workers in 2011 (in M)	27.9	100.0%	8.9	32.0%	16.1	57.6%	2.9	10.4%
# workers in 2016 (in M)	31.5	100.0%	10.0	31.7%	18.3	58.2%	3.2	10.2%
Observations	1,660		525		963		172	

Notes: Table shows mean statistics weighted with firm stratification weights. In Panel A and Panel B uneven columns give the capital stock in 1,000 € and even columns give the relative shares of the capital stocks across technology classes in % adding up to 100 per column. In Panel C uneven columns give the number of firms or workers employed in M and even columns give the relative shares of firms or workers across non-adopting and adopting firms in % adding up to 100 per row. Note that the number of firms is time-constant due to the balanced panel of firms, see section 2.2.

Table 2 also provides information on the technology usage across all firms and by adoption group in 2011 (Panel A) and 2016 (Panel B) both in absolute levels (uneven columns) and percentage shares (even columns).[18] The average German firm across all adoption groups mainly invested in modern technologies since both the share of digital and frontier technologies increased at the expense of a decrease in the relative importance of older technologies (see columns 1-2). The average German firm in 2011 mostly employed older manual technologies, but predominantly used digital technologies in 2016. Frontier technologies play a small but growing role in the German economy with more than 10% of the average firm's capital stock being based on frontier technologies in 2016.[19]

Columns 3-8 document distinct patterns of technology use by adoption group. Panel A

[18] Note that these are not identical to the technology-specific capital shares k_{fgt} from the survey. For comparison, these can be found in Appendix Table A.1.

[19] Appendix Table A.2 shows that these average technology patterns are similar when differentiating between manufacturing and service firms.

shows that the initial capital stock of non-adopting firms predominantly consist of older manual technology (column 3), which accounts for more than two-thirds of their technology (column 4). Moreover, investments until 2016 left the technology distribution unchanged. In contrast, frontier and digital adopters predominantly used digital technologies in 2011 (columns 5 and 7). With respect to the initial use of and investments into frontier technologies, digital and frontier adopters clearly differ: while for digital adopters frontier technologies in 2011 only account for 5% of the capital stock (Panel A, column 6), they already represent more than 20% among frontier adopters (Panel A, columns 8). Both adopter groups extended the share of frontier technologies until 2016, although frontier adopters invested much more in frontier compared to digital technologies, in absolute terms. In 2016, frontier technologies represent 40% of the capital stock of the average frontier adopter (Panel B, column 8), while the corresponding share remains rather low with 7.6% for digital adopters (Panel B, column 6).[20] Hence, the adopter groups resulting from our classification differ significantly in their capital structure.[21] In the remainder of the paper, we use these adoption groups to examine whether and how these adoption groups contribute to changes in the German occupational composition, but first look at the factors determining a firm's investment choice.

3.2 Determinants of firms' technology adoption

Our novel measure of firm-level technology adoption allows differentiating between computer-controlled digital technologies and frontier technologies and our survey covers firms from both manufacturing and service industries. Therefore, we are the first to provide insights on factors underlying a firm's decision to invest in *frontier* rather than computer-controlled digital technologies. For this, we investigate which firm-specific initial period characteristics influence this important decision margin by estimating the following regression specification:

$$P(g\text{-adopter}) = \alpha + \beta X_{f,t=2011} + \epsilon_f \tag{4}$$

where $P(g\text{-adopter})$ denotes either the probability of a firm f to mainly adopt any type of modern technology $g \in (3.0; 4.0)$ (Table 3, columns 1-3) or the probability of a firm to mainly adopt frontier rather than computer-controlled digital technologies (Table 3, columns 4-6) and $X_{f,t=2011}$

[20]In line with our study, Acemoglu et al. (2022a) find that the adoption of AI, robotics, dedicated equipment, specialized software, and cloud computing remains low across firms in the U.S.

[21]Appendix Table A.3 also shows differences in initial firm and workforce characteristics by adoption group.

is a set of initial-period firm and workforce characteristics which we choose based on the main potential forces behind a firm's decision to invest in modern technologies in line with the existing theoretical and empirical literature. To estimate the probability to mainly adopt any type of modern technologies $g \in (3.0; 4.0)$, we use as a dependent variable in columns 1-3 an indicator that is unity if a firm is either a digital or frontier adopter and zero otherwise. In columns 4-6, we exclude non-adopting firms and use an indicator as the dependent variable which is unity if firm f is a frontier adopter and zero for digital adopters.[22]

Firm's initial size and occupational composition determine modern technology adoption. Starting with modern technology adoption in either digital or frontier technologies, columns 1-3 in Table 3 document that manufacturing firms are about 25 percentage points less likely to adopt new technology than service firms. This highlights the importance of including service industries in the scope of the analysis when examining the adoption of new technology. Technology adopting firms in our sample are ex-ante larger than non-adopters (column 1), which is in line with previous studies documenting that firms adopting modern technologies are larger since firm size facilitates the amortization of investments (Koch et al., 2021; Bonfiglioli et al., 2020; Acemoglu et al., 2020; Zolas et al., 2020).[23] With respect to the workforce composition, we find that firms' share of routine workers positively correlates with the probability of adopting any modern technology (columns 1-3). This aligns with the finding of Koch et al. (2021) that firms with a high share of production tasks are more likely to adopt robots as they may have incentives to use this technology to replace these tasks. More generally, previous literature has shown that modern automation technologies substitute for routine jobs that follow well-defined rules (among others, e.g., Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2018). Thus, the positive correlation between the initial share of routine workers and the propensity to adopt modern technology we find in Table 3 might point to the potential labor-capital substitution potential in adopting firms. Additionally, the positive correlations between the employment share of experts and specialists and the probability of a firm adopting any type of modern technology (columns 1-3), suggest that firms also need skills that are complementary to implementing and

[22]Note that by estimating the investment decision this way, we assume independence of irrelevant alternatives between the three adoption choices. Relaxing this assumption by estimating a multinomial probit, yields qualitatively similar results, see Appendix Table A.4. Hence, we prefer the simpler specification.

[23]Similarly, Zolas et al. (2020) rely on a large-scale firm survey in the U.S. and find that AI adoption is rare and generally skewed towards larger firms. Equally, Acemoglu et al. (2022a) find for the U.S. that the adoption of AI, robotics, dedicated equipment, specialized software, and cloud computing remains low (especially for AI and robotics), varies substantially across industries, and concentrates on large and young firms.

Table 3: Determinants of firm-level technology adoption

	Adopting any modern technology			Adopting frontier technology		
	(1)	(2)	(3)	(4)	(5)	(6)
Manufacturing firm	-0.261***	-0.252***	-0.261***	-0.087*	-0.070*	-0.077*
	(0.063)	(0.063)	(0.063)	(0.046)	(0.042)	(0.043)
Firm wage premium		0.003	0.000		-0.003	-0.004
		(0.005)	(0.005)		(0.003)	(0.003)
Firm employment	0.040**		0.044**	-0.017		0.014
	(0.016)		(0.018)	(0.017)		(0.016)
MDF port within reach	-0.078	-0.072	-0.084	0.138**	0.134**	0.130**
	(0.095)	(0.095)	(0.095)	(0.058)	(0.061)	(0.061)
Worker shares by main task of occupation (reference: NRM)						
NRC	0.267***	0.253***	0.262***	0.133	0.127	0.133
	(0.092)	(0.093)	(0.092)	(0.085)	(0.085)	(0.085)
RC/RM	0.311***	0.292***	0.308***	0.061	0.050	0.059
	(0.072)	(0.073)	(0.072)	(0.073)	(0.069)	(0.071)
Worker shares by required skill level (reference: Professionals)						
Experts	0.330***	0.340***	0.328***	-0.035	-0.043	-0.043
	(0.107)	(0.106)	(0.108)	(0.120)	(0.111)	(0.111)
Specialists	0.223**	0.231**	0.231**	0.020	0.098	0.103
	(0.104)	(0.106)	(0.107)	(0.121)	(0.115)	(0.115)
Unskilled workers	-0.194	-0.187	-0.198	0.282	0.242	0.234
	(0.127)	(0.125)	(0.127)	(0.200)	(0.188)	(0.191)
Worker shares by age group (reference: 40-55 years of age)						
<40 years of age	-0.123	-0.124	-0.114	-0.169*	-0.141	-0.142
	(0.087)	(0.088)	(0.088)	(0.096)	(0.089)	(0.089)
>55 years of age	0.106	0.118	0.122	-0.323**	-0.199*	-0.193*
	(0.149)	(0.146)	(0.152)	(0.132)	(0.109)	(0.112)
Constant	0.399***	0.437***	0.381***	0.061	0.003	-0.016
	(0.117)	(0.125)	(0.125)	(0.083)	(0.077)	(0.083)
Observations	1,660	1,660	1,660	1,135	1,135	1,135
R-squared	0.155	0.147	0.155	0.082	0.119	0.120

Notes: Columns 1-3 are based on the entire firm sample, while columns 4-6 are restricted to adopting firms. All models are estimated with firm stratification weights. The variable manufacturing firm stems from a self-assessment during the firm survey of whether the firm should be considered as a manufacturer or service provider. Firm employment enters in logarithmic form. Firm wage premium corresponds to firm fixed effects for the period 2003-2010 based on the method by Abowd et al. (1999). Includes dummy for missing high-speed internet access information (10 firms) and mostly small firms with missing firm wage premium (124 firms). NRC: non-routine cognitive jobs; RC/RM: routine cognitive/manual jobs; NRM: non-routine manual jobs. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

maintaining modern technologies. The positive correlation between the share of NRC jobs and the probability of adopting modern technology additionally supports this hypothesis. This is in line with previous research findings that firms need high-skilled employees conducting non-routine analytical tasks to plan, implement, and maintain modern technologies (Harrigan et al., 2021; Bresnahan et al., 2002).

Frontier technology adoption is mainly driven by high-speed internet access. The decision to adopt the most recent frontier technologies rather than computer-controlled digital technologies (columns 4-6) shows only a few significant correlates. In particular, manufacturing firms and firms with a high share of older workers are less likely to mainly adopt frontier technologies. This finding is in line with studies showing a negative relationship between firms' workforce age and the probability of technology or software adoption (Meyer, 2011). The availability of a high-speed internet access point turns out to be the single most important correlate for the decision to mainly invest in frontier technologies.[24] Since frontier technologies connect physical and digital spheres to perform work processes automatically without human intervention, these technologies are fully integrated into the firms' central IT systems and require high-speed internet. This insight complements previous studies emphasizing the role of firm's ICT infrastructure for technology adoption, highlighting that the diffusion of high-speed broadband internet correlates positively with the adoption of computers (DeStefano et al., 2018; Nicoletti et al., 2020). Thus, the decision to adopt frontier rather than digital technologies significantly relates to conducive conditions such as a favorable age structure and high-speed internet access, but cannot be explained by initial differences in firm size and workforce composition. This might suggest that the adoption of frontier technologies to a larger degree depends on having the right IT infrastructure as well as young (and technology-open) workers to implement the new technology.

For our subsequent analysis of changes in the occupation structure in Germany, these results point to significant initial differences in size and occupation structure between adopters and non-adopters and nuanced differences between digital adopters and frontier adopters. Initial period heterogeneity between adoption groups may also induce aggregate employment shifts across occupation groups, which we explore in the next section.

4 De-routinization in the German economy

In this section, we aim to understand how the adoption of frontier technologies contributes to changes in the occupational composition of the German economy. We start by examining

[24]We control for the availability of high-speed internet by including an indicator of whether a firm's location is within the reach of a high-speed internet access point. Beyond these access points, internet speed deteriorates markedly. Falck et al. (2014) show that the roll-out of the first generation of digital subscriber line (DSL) technology in Germany was entirely built on the layout of the pre-existing voice telephony network to minimize costs. Since the telecommunication data exploits historical peculiarities, the exogenous variation in the DSL roll-out can be used to exogenously estimate the effects of high-speed internet availability.

occupational shifts in the economy, before assessing the subsequent contributions of non-adopters, digital adopters, and frontier adopters.

4.1 Aggregate de-routinization in Germany

Aggregate decline in routine jobs between 2011 and 2016. Overall, Figure 1(a) documents a substantial decline in routine cognitive and routine manual (RC/RM) jobs by about 2.4 percentage points among the German workforce from 2011 to 2016. This decline is mostly compensated by rising non-routine cognitive (NRC) employment shares (+1.8pp) and to a smaller extent by rising non-routine manual (NRM) employment shares (+0.6pp). These shifts are in line with the literature that suggests that technology adoption aims to substitute for human labor in routine, replaceable tasks, while complementing workers in non-routine, non-replaceable tasks (Acemoglu and Autor, 2011).

Figure 1: De-routinization in the aggregate and by adoption group between 2011 and 2016

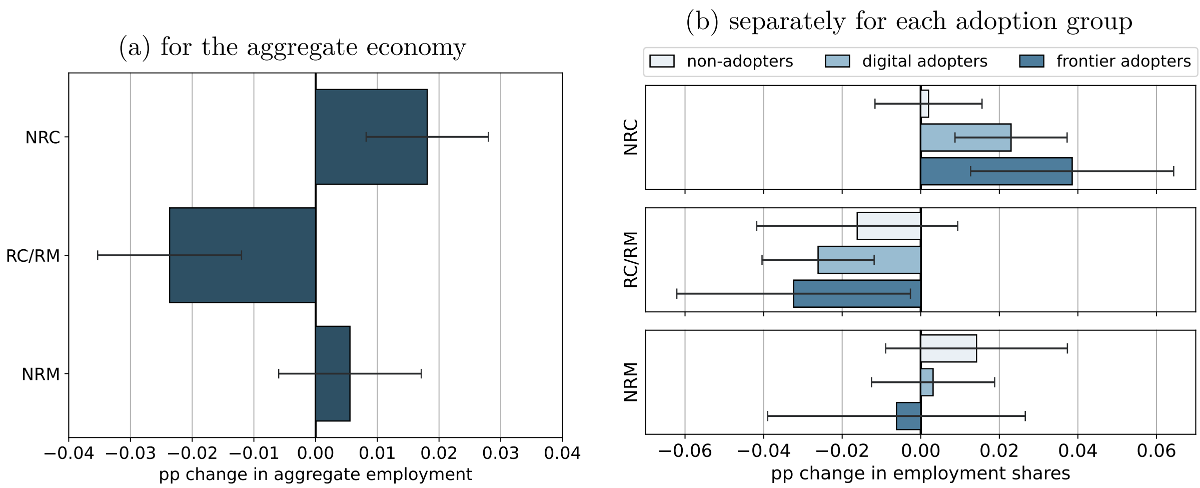

Notes: Panel (a) shows aggregate occupational shifts by task domain for the entire German workforce; Panel (b) depicts the change in the share of occupations by task domain among the workforce employed at adoption group g. Weighted with time-varying employment-weighted firm stratification weights $\tilde{s}_f$. NRC: non-routine cognitive jobs; RC/RM: routine cognitive/manual jobs; NRM: non-routine manual jobs. Confidence bands based on jackknife standard errors.

Routine jobs decrease most in the workforce of frontier adopters. The key advantage of our data is that we can distinguish those shifts by technology type (non-adopters, digital adopters, and frontier adopters) and trace occupational employment shifts at the firm level. Hence, Figure 1(b) shows the employment de-routinization separately for each adoption group. More precisely, it shows the change between 2011 and 2016 in the share of workers that belong to a certain task group and are employed at the respective adoption group, for example, frontier adopters. The proportion of routine jobs decreases the most in the workforce of frontier adopters, followed by a somewhat smaller decrease in the proportion of routine jobs in the workforce of digital adopters. Moreover, de-routinization among both adopter groups goes along with increasing shares of workers employed in non-routine cognitive jobs. Figure 1(b) provides first empirical firm-level evidence that the adoption of frontier technologies hence continue to contribute to and even slightly accelerate aggregate de-routinization, rather than showing a totally different employment shift as expected by some scholars (Webb, 2020; Brynjolfsson et al., 2018).[25]

Interestingly, de-routinization occurs among all adoption groups including non-adopters, though not statistically significant. This indicates that de-routinization is partly driven by forces other than technological change, such as domestic outsourcing (Cortes and Salvatori, 2019), offshoring (Reijnders and de Vries, 2018), or changes in consumer demand (Autor and Dorn, 2013). Contrary to adopters, de-routinization among non-adopters favors non-routine manual rather than non-routine cognitive jobs.

4.2 Economic reasoning about the sources of de-routinization

The observed aggregate decline in routine employment suggests that frontier technologies continue to contribute to recent de-routinization. Yet, three potential channels may contribute to this outcome, and understanding the relevance of each has important implications:

Channel 1: Direct routine replacement within technology-adopting firms. The RRTC literature builds on the idea that computer-controlled technologies induce shifts *within firms* due to capital-labor substitution, reducing the shares of routine, unskilled, and low-paid workers within firms. Several studies find empirical evidence that firm-level technology investments in ICT technology and industrial robots are associated with or induce labor demand shifts towards non-routine, more skilled, and better-paid workers who are complementary to the new technologies (Gaggl and Wright, 2017; Böckerman et al., 2019; Acemoglu et al., 2020; Bonfiglioli et al., 2020; Humlum, 2019; Koch et al., 2021). Yet, it remains an open question whether the rapidly emerging frontier technologies associated with the fourth industrial revolution lead

[25] In an additional analysis, we show that these occupational employment shifts are accompanied by shifts between skill groups (Appendix Figure D.5) and wage groups (Appendix Figure D.6). In particular, we find a clear trend towards skill and wage polarization: Both the highest and the lowest skill group (wage group) increase at the expense of the medium skill group (wage group). However, the increase in experts and high-paid jobs group is solely driven by frontier adopters while the increase in helpers and low-paid jobs group is also driven by non-adopters and digital adopters.

to job displacement within adopting firms and which occupation group will be most exposed (Brynjolfsson et al., 2018; Webb, 2020).[26] If the average employment shift observed within frontier adopters – the *within firms* effect - resembled the pattern in Figure 1(b), this would suggest a widespread replacement of routine jobs within adopting firms. In this case, frontier technologies would be routine replacing per se, i.e., independent of any firm-specific conditions.

Channel 2: Employment reallocation between adoption groups. The RRTC literature discusses a second mechanism that may induce employment shifts between firms: initial heterogeneity in firm size and employment structure between technology-adopting and non-adopting firms. This mechanism builds on the idea that technology adoption allows adopting firms to grow faster which results in aggregate occupation shifts if adopters and non-adopters differ substantially in their workforce composition. As we show in section 3.2, adopters differ from non-adopters in their initial characteristics: adopters are larger, more productive, and have initially higher shares of replaceable routine jobs. Due to technology adoption, these firms might grow faster. This would imply a *between adoption groups* effect: Technology-adopting firms differ in their initial employment structure from non-adopting firms and grow faster than the latter. This might induce a shift in overall employment towards the group of adopters that have high shares of routine jobs, resulting in an aggregate shift towards routine jobs. Harrigan and Reshef (2015) similarly provide a theoretical model that combines firm heterogeneity in productivity with firm heterogeneity in skill structures. A strong between adoption groups effect would point towards the heterogeneity between adopters and non-adopters and the role of an unequal IT infrastructure for implementing frontier technologies.

Channel 3: Employment reallocation within adoption groups. As a third channel, the intensity of de-routinization induced by technology adoption might vary between firms within the group of frontier adopters. If frontier adopters are heterogeneous with respect to the initial conditions when adopting new technologies, we might find stronger substitution in places where related adjustments are more beneficial and easier to implement. For instance, technology adoption requires complementary investments into training and organizational change

[26] Acemoglu et al. (2022b) link AI exposure to changes in skills demanded by establishments in job postings, finding that AI technology alters the task structure of jobs by generating new tasks accompanied by new skill demands. Acemoglu et al. (2022a) find suggestive evidence for skill upgrading by firms using advanced technologies, though their data does not allow to quantify the contribution of technology adoption at the firm level to occupational shifts. We complement these recent studies by linking the firm-level adoption of frontier technologies to occupational employment shifts in the workforce and thereby quantifying their contribution to de-routinization at the micro- and macro-level.

(Brynjolfsson and Hitt, 2000; Bresnahan et al., 2002; Brynjolfsson and Milgrom, 2013), as it is usually not sufficient to buy off-the-shelf technologies and adopt it in the previous environment without any accompanying adjustments. Firms may be differently equipped for such investments and restructuring. This can induce a *within adoption groups* effect, i.e., a shift in aggregate employment shares that occurs due to compositional changes between technology-adopting firms. In fact, Cortes et al. (2023) combine a Melitz (2003) type model with skills and task-biased technological change (TBTC) and find that task changes differ in intensity also within the group of adopters depending on their size and productivity. This view is also in line with Battisti et al. (2023) who find that the combination of technologies with organizational change is associated with changes in tasks.

This heterogeneity between adopter firms could itself be either due to *scale effects* or due to *composition effects*. Scale effects occur when the change in the occupation structure differs between firms that differ in initial firm size – for example larger declines of routine jobs in larger firms. Composition effects occur when firms with different initial occupational employment structures experience differences in employment growth – for example, faster growth among non-routine-intensive firms.

A pronounced within adoption groups effect would thus point to specific conditions that are necessary for frontier technologies to be routine replacing. Hence, a within adoption groups effect could reveal the conditions under which frontier technologies are profoundly routine replacing and under which conditions the adoption does not or only moderately replaces routine work.

Which of the three main channels (and sub-channels) contributes most to de-routinization is an important empirical question. Whether the strength of de-routinization from the adoption of frontier technologies follows its inherent nature (within firms effect), results from unequal access to adopting these technologies (between adoption groups effect), or depends on specific conditions between adopting firms (within adoption groups effect) comes with very different policy implications. To quantify the relative contribution of each channel, we perform a decomposition analysis in the next section.

4.3 Decomposing aggregate de-routinization

In order to decompose aggregate de-routinization, we build on Autor et al. (2020) and Acemoglu et al. (2020) and decompose aggregate occupational employment shifts into those that occur *within firms*, between firms and *between adoption groups*, and between firms but *within adoption* *groups*.[27] In particular, we decompose the change in aggregate employment share (e.g., routine jobs), $\Delta\lambda$ as follows:

$$\Delta\lambda = \underbrace{\sum_g \Delta\bar{\lambda}_g}_{\text{within firms}} + \underbrace{\sum_g \Delta\left(\bar{\lambda}_g - \bar{\lambda}\right)\left(\tilde{s}_g - \frac{G}{F}\right)}_{\text{between adoption groups}} + \underbrace{\sum_g \Delta \sum_{f \in g}\left(\lambda_f - \bar{\lambda}_g\right)\left(\tilde{s}_f - \frac{w_f}{F}\right)}_{\text{within adoption groups}} \tag{5}$$

where λ is the aggregate employment share of an occupation, $\Delta\lambda$ is its change over time, $\bar{\lambda}_g$ is the average firm-weighted share of an occupation in adopter group g, $\bar{\lambda}$ is the average firm-weighted share of an occupation in the aggregate economy, $\tilde{s}_g$ is the share of adopter group g in total employment, G and F are the weighted total number of firms of adopter group g and across all groups, respectively, λ_f is the share of an occupation in firm f, $\tilde{s}_f$ is the employment share of firm f, and w_f is the survey weight of firm f. The *first* component of equation (5) captures the average occupation change within firms over time by adopter group, $\bar{\lambda}_g$, for $g \in \{1.0/2.0(non); 3.0(digital); 4.0(frontier)\}$. The *second* component captures changes in deviations of the group-specific mean from the average employment change in the economy, $\bar{\lambda}$. To restore representativeness regarding total employment, this term is weighted by deviations between the group-specific employment share, $\tilde{s}_g$, and the relative number of firms of each adopter group (G/F). The *third* component captures deviations in firms' occupation change from the group-specific mean. This term is again weighted by the deviations of firms' employment shares, $\tilde{s}_f$, from firms' relative survey weight (w_f/F). Each component of equation (5) is calculated for each adopter group. Summing each component over the three adopter groups and adding up the three components yields the total aggregate occupation change in the economy. The details of the decomposition are shown in Appendix C.

Based on equation (5), we conduct the decomposition for each adopter group and construct confidence bands based on jackknife standard errors. In the following discussion, we focus on the results for the group of frontier adopters, which experience the fastest de-routinization. Additional, mostly insignificant results for digital adopters and non-adopters can be found in Appendix D. Moreover, Appendix E provides extensive robustness checks for our decomposition analysis.

[27] In contrast to classic within-between decompositions such as, e.g., Harrigan et al. (2021), we further decompose aggregate changes by adopter groups and further differentiate between-firm shifts into those that take place within adopter groups from those occurring between adopter groups. Also note that our within-firms component is free of any firm-size related weights and therefore represents an average within-firm shift unrelated to firm size.

Sources of de-routinization for firms adopting frontier technology. Figure 2 presents the decomposition of the de-routinization among frontier adopters into the three components of equation (5). It shows that the aggregate decline in routine jobs among frontier adopters is not driven by within firms effects (first component in equation 5). In fact, the average frontier adopter (insignificantly) increases rather than decreases the share of routine jobs at the firm level. Hence, de-routinization among frontier technology adopters is not a simple story of new technologies generally substituting for routine jobs, but must be driven by between-firm shifts. This result is in line with Harrigan et al. (2021) who document that the aggregate rise of technology-related occupations in France between 1994 and 2007 is not driven by within-firm adjustments but by changes in the composition between firms. However, complementing this evidence, we find that between firm shifts is not taking place between adoption groups. The between adoption groups component (second component in equation 5) turns out to be a negligible source of the contribution of frontier adopters to aggregate employment changes.

Instead, the contribution of frontier technology adopters to the aggregate de-routinization is explained by differences between firms within the group of frontier adopters, i.e., the within adoption groups component (third component in equation 5). Accordingly, heterogeneity between firms *within* the group of frontier adopters explains the contribution of frontier adopters to aggregate de-routinization. Put differently, only a subset of firms with specific initial conditions contribute to de-routinization. We next aim to better understand this heterogeneity by further decomposing the within-group component.[28]

Zooming into the within adoption groups effect. To delve into mechanisms behind the within adoption groups component, Figure 3 further decomposes the within adoption groups effect into a scale, composition, and simultaneous effect. The scale effect reflects differences in occupational employment changes by firms that differ in initial firm size. The composition effect reflects differences in firm growth between firms that differ in initial occupational structures, and the simultaneous effect captures the combination of the two.[29] Figure 3 highlights that both, scale and composition effects play a sizeable role in the decline in routine jobs and the rise of

[28]Appendix D documents that digital technology adopters contribute to the aggregate de-routinization by shifts within firms and shifts between firms within the group of digital adopters, both components are equally in magnitude but statistically insignificant. These results are in line with Böckerman et al. (2019) documenting that firm-level ICT usage leads to de-routinization within adopting firms. Also note the absence of any significant sub-components of the within adoption groups component for digital adopters. For non-adopters, neither the overall contribution to de-routinization is significant, nor any of the sub-components.

[29]The respective decomposition formula and some further explanation can be found in equation (13) in Appendix C.

Figure 2: Sources of de-routinization between 2011 and 2016 for frontier technology adopters

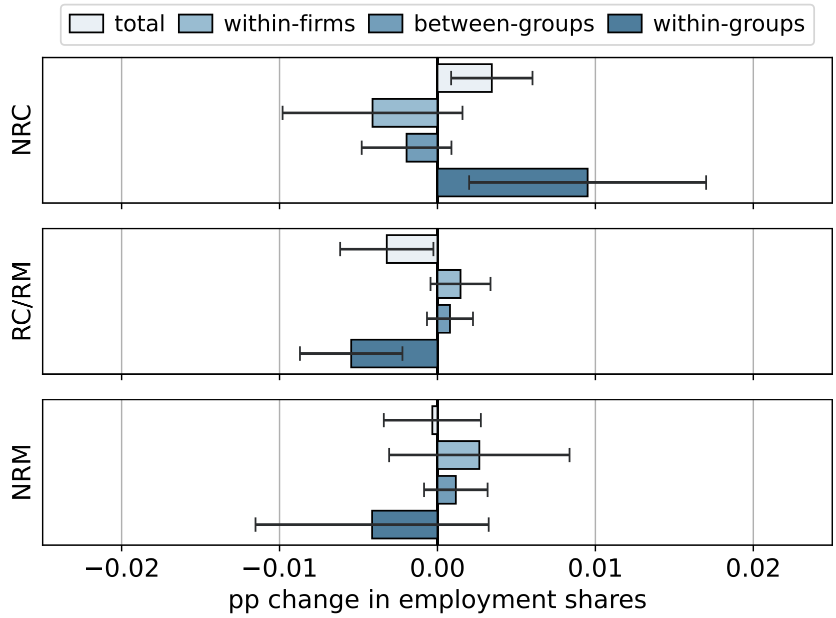

Notes: The Figure reports the contribution to aggregate de-routinization between 2011 and 2016 by frontier adopters differentiating between the three components from equation 5: the contribution of the within firms, between adoption groups and within adoption group effects. Weighted with firm stratification weights. NRC: non-routine cognitive jobs; RC/RM: routine cognitive/manual jobs; NRM: non-routine manual jobs. Confidence bands based on jackknife standard errors.

non-routine cognitive jobs, while the simultaneous effect is negligible.[30] Hence, the aggregate decline in the routine employment share among the workers employed at frontier adopters occurs because (1) initially larger frontier adopters reduce their routine employment share in contrast to smaller frontier adopters (scale effect) and because (2) initially less routine-intensive firms experience stronger employment growth (composition effect).

Taken together, the findings from the decomposition suggest that the adoption of frontier technology does not lead to a widespread replacement of routine jobs within firms. Rather, only a subset of larger firms is responsible for the de-routinization among frontier adopters. At the same time, firms that already have a higher share of non-routine cognitive jobs are better able to gain benefits from technology adoption and thus grow faster. In the next section, we investigate the source of the scale and composition effect to gain a better understanding of the underlying mechanisms.

[30]Appendix Figure D.1 shows the corresponding decomposition of the between adoption groups effect of frontier adopters into size, composition, and simultaneous effects, none of which are significant.

Figure 3: Sources of the within adoption groups effect for frontier technology adopters

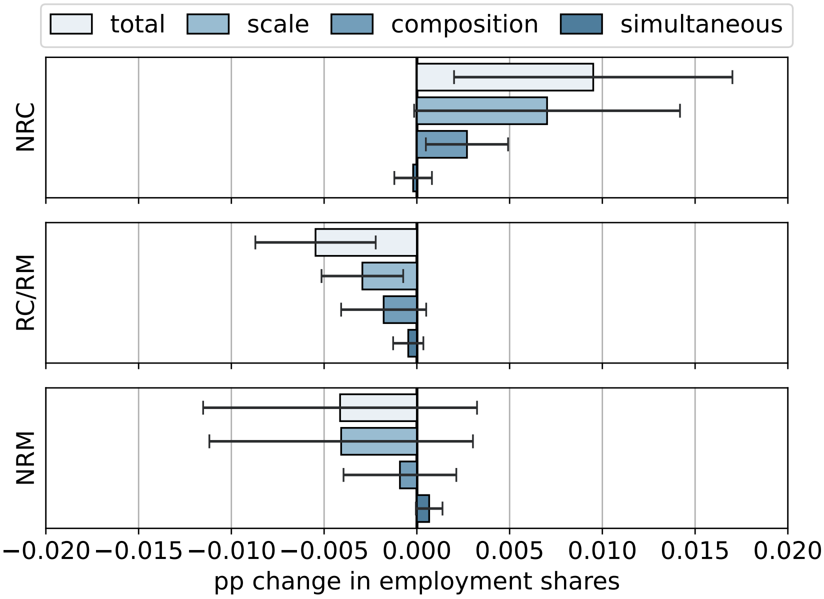

Notes: The Figure reports the decomposition of the within adoption groups effect for frontier adopters as shown in Appendix equation (13): contribution of frontier adopters to aggregate de-routinization between 2011 and 2016 due to occupational employment shifts induced between frontier adopting firms related to size, composition differences, and a combination of both. Scale, composition and simultaneous components add up to the total within adoption groups effect of frontier adopters. NRC: non-routine cognitive jobs; RC/RM: routine cognitive/manual jobs; NRM: non-routine manual jobs. Confidence bands based on jackknife standard errors.

5 Delving into the heterogeneity among frontier adopters

In this section, we aim to understand whether de-routinization among frontier adopters is driven by genuine scale and composition effects, or whether these effects rather reflect other factors that the decomposition framework does not incorporate. For this, we turn to a regression approach to estimate the roles of initial occupational composition and firm size for de-routinization, while controlling for various potentially related firm characteristics.[31]

5.1 Scale effect among frontier adopters

While the decomposition revealed a size-related decline in routine jobs among frontier adopters (scale effect), these employment shifts might reflect other characteristics related to larger frontier adopters, such as the sector of activity, the firm's age, or a workforce composition favoring routine replacement. To explore the relevance of such characteristics, we regress changes in occupational employment shares on firm size and subsequently extend the set of control variables. If these additional factors eliminate the size effect, this gives us insights into potential mechanisms

[31] An alternative to the regression analysis would be a decomposition within more homogeneous subgroups. However, our sample size of 172 firms adopting frontier technologies makes such a further sample split infeasible.

underlying the size effect. In particular, we estimate:

$$\Delta\lambda_f = \beta_0 + \sum_{i=2}^{4} \gamma_i S_{if} + \beta\, X_f + \epsilon_f \tag{6}$$

where $\Delta\lambda_f$ is the change in firm f's employment share of a given occupation group, S_{if} is a set of firm-size dummies where firms with fewer than 10 workers are the reference category, and X_f is a set of further controls that potentially explain the heterogeneity across firms.

Table 4 shows the results.[32] Consistent with the decomposition, regressing the change in RC/RM jobs on firm size without further controls confirms the scale effect, i.e., significant negative estimates increasing in magnitude for the firm-size dummies (column 1). However, this scale effect might be driven by size-related heterogeneity in core firm characteristics such as sector of activity, firm age, or general firm performance.[33] Yet, controlling for these characteristics in column 2 does not weaken, but rather strengthens the scale effect. The change in routine employment is 12.9pp lower for large firms compared to small firms, which is sizable compared to the average shift of 1.85pp.

The scale effect might also reflect size-related differences in the initial workforce composition. For instance, larger firms might have initially higher shares of routine workers and therefore experience stronger routine-replacing effects. Adding initial employment shares across the main task type (3 categories), skill shares (4 categories), age shares (3 categories), and the share of IT workers in column 3, however, does not reduce the scale effect. Thus, the firm's initial workforce structure does not contribute to the observed size-related heterogeneity in the task shifts among frontier adopters.

Another size-related source of heterogeneity might stem from capital investments. Although all firms in our regression analysis are frontier adopters, there might be heterogeneity in the initial technology-specific capital stock in 2011 or technology-specific investments between 2011 and 2016. For example, larger firms may already have been better equipped with frontier technologies in 2011 or may have invested more in frontier technologies since 2011 than smaller frontier adopters. Yet, column 4 speaks against this hypothesis. Thus, initial technology use and even the quantity of investments also do not explain the size-related heterogeneity in task shifts among

[32]Note that Table 4 displays the result for the change in RC/RM employment since the NRC results are almost a mirror image of the RC/RM decline and changes in RC/RM, NRC and NRM employment add up to zero.

[33]We capture firm performance by the percentile in which the firm belongs across the wage premia, estimated by firm fixed effects based on the approach proposed by Abowd et al. (1999).

Table 4: Decline in routine employment of frontier adopters between 2011 and 2016

Dependent variable: Change in firm-level employment share of routine jobs $(\Delta RC/RM)$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Firm size (reference: <10 employees):							
10-49 employees	-0.011	-0.081*	-0.111	-0.070	-0.042	-0.060	-0.007
	(0.037)	(0.048)	(0.067)	(0.061)	(0.045)	(0.050)	(0.042)
50-199 employees	-0.049**	-0.131***	-0.158***	-0.137**	-0.009	-0.073	0.004
	(0.019)	(0.031)	(0.044)	(0.054)	(0.047)	(0.056)	(0.058)
200+ employees	-0.058***	-0.129***	-0.145***	-0.192***	-0.098*	-0.115	-0.062
	(0.018)	(0.028)	(0.042)	(0.055)	(0.052)	(0.071)	(0.072)
Adj R^2	0.004	0.094	0.160	0.290	0.591	0.447	0.536
Firm characteristics		✓	✓	✓	✓	✓	✓
Workforce structure			✓	✓	✓	✓	✓
Capital structure				✓	✓	✓	✓
Perceptions					✓		
Skill demands						✓	✓
Training							✓

Notes: $N = 172$. The average firm-weighted shift in routine jobs among the 172 frontier adopters is +1.85pp (note that this corresponds to the respective figure in Figure 2 divided by 0.079 which is the share of frontier adopters among all firms, see Table 2). The set of firm characteristics include a dummy for the self-assessment of whether the firm should be considered as a manufacturer or service provider, sector (5 categories), firm age (dummy if founded after 2001), firm performance (AKM fixed effects). The workforce structure characteristics refer to the initial values in 2011 and include initial employment shares across main task types (3 categories), skill shares (4 categories), age shares (3 categories), and the share of IT workers. The capital structure characteristics include the initial level of technology-specific capital stocks and the change in the technology-specific capital stocks between 2011 and 2016. Perceptions include the assessment of chances and risks related to the adoption of frontier technologies (16 dimensions), see Appendix F.1. Skill demands include the assessment of how skill demands changed in the past five years (14 dimensions), see Appendix F.2. Training includes a set of variables describing how the extent and content of vocational and further training changed in the past five years (10 dimensions), see Appendix F.3.

frontier adopters.[34]

One remaining explanation for the scale effect might be that firm size affects the type of frontier technology that is implemented or the way how such technologies are implemented. To get tentative evidence on the role of such factors, column 5 adds subjective perceptions from the survey about the chances and risks associated with the use of frontier technologies. In total, we asked the interviewees to assess how they perceive potential chances and risks of adopting frontier technologies. The 16 items include potential risks such as high investment costs, restructuring needs, and labor shortages, but also potential benefits such as labor cost savings. See Appendix F.1 for more information about all 16 perception dimensions.[35] Column 5

[34] The size effect does also not diminish when controlling for firm differences in churning and turnover rates (pre-2011) as an indicator of a firm's ability to quickly restructure its workforce, or in entry wages as an indicator of the firm's ability to hire new NRC workers.

[35] Note that firms were asked a random subset of perceptions in the survey. We impute the randomly missing

of Table 4 shows that the inclusion of these perceptions substantially lowers the coefficient on the firm-size dummies and eliminates statistical significance except for the largest group of firms.

The results so far tentatively indicate that the size heterogeneity in occupational employment shifts partially reflects that larger firms differ in why they invest in frontier technologies and how they think about implementing them. One particular aspect of implementing frontier technologies could be the adjustment to changing skill requirements. If larger firms invest in frontier technologies that change skill requirements more profoundly, this may necessitate stronger adjustments of the workforce to changing skill demands. To test this channel, column (6) alternatively adds indicators from the survey about how firms view changing skill requirements during the last five years with respect to 14 skill dimensions.[36] This results in a large drop of coefficients compared to column 4 and entirely eliminates the statistical significance of the size effects.

Changing skill requirements should be accompanied by changing training needs. Adding survey information on how firms assess how vocational and further training needs and content changed between 2011 and 2016 in column 7, further reduces the coefficient sizes for the size dummies.[37]

Overall, these results highlight that the scale effect is remarkably stable to a large set of controls as it cannot be explained by any initial differences in firm workforce characteristics or by prior investments in frontier technologies. Our findings tentatively suggest that the size heterogeneity reflects differences in which frontier technologies get adopted and how they get implemented. The heterogeneity between large and small frontier adopters in implementing frontier technologies is also reflected in how they perceive changes in skill requirements and training needs. Hence, the adoption of frontier technologies in larger firms boosts de-routinization compared to smaller firms.

---

items as described in Appendix F.1. Using the original items and a missing indicator for randomly missing items provides comparable results. As a drawback of the randomization, it is infeasible to isolate the marginal effect of each perception measure. Hence, we only discuss their joint effects.

[36]See Appendix F.2 for more information about the 14 skill dimensions. Note that firms were asked a random subset of skill demands in the survey. We impute the randomly missing items as described in Appendix F.2. Using the original items and a missing indicator for randomly missing items provides comparable results.

[37]See Appendix F.3 for more information about the ten items on how vocational training and further training changed in the past five years. We impute missing items as described in Appendix F.3. Using the original items and a missing indicator for missing items provides comparable results.

5.2 Composition effect among frontier adopters

We follow a similar approach to shed light on why frontier adopters with an initially high share of NRC employment (low share of RC/RM employment) grow faster compared to the average frontier adopter. For this, we regress firms' employment growth on firms' initial occupational employment shares and subsequently extend the set of control variables:

$$\Delta \log N_f = \beta_0 + \gamma_i \lambda_{if} + \beta X_f + \epsilon_f \tag{7}$$

where $\Delta \log N_f$ gives the log employment growth between 2011 and 2016 and λ_{if} gives the initial firm-level share of a given occupation group $i \in \{NRC; RC/RM; NRM\}$ for firm f with routine jobs being the reference category.

Table 5 shows the results. Without any further controls, column 1 confirms the significantly positive correlation between the share of NRC employment and firm employment growth from the decomposition. Adding firm characteristics related to the sector of activity, firm age, and firm performance in column 2, does not weaken this positive growth effect, i.e., these characteristics do not explain why a high initial NRC share results in stronger employment growth among frontier adopters. This growth effect also remains stable when adding the firm's workforce structure in terms of age and IT personnel in column 3.[38] In column 4, we test whether composition-related heterogeneity stems from differences in the capital structure with respect to the initial capital stock or technology-specific investments between 2011 and 2016. Yet, column 4 shows that adding these indicators does not affect the strength or direction of the NRC effect.

Contrary to the scale effect in the previous subsection, neither adding the perceived chances and risks of frontier technology adoption (column 5), the changing skill demands (column 6), nor adding the changes in vocational and further training needs (column 7) mitigate the composition effect.[39] In fact, the growth-enhancing effect of an initially high NRC-intensive workforce becomes slightly stronger when adding these controls. Hence, firms with an initially higher NRC employment share are better prepared to benefit from the growth potential of a given set of frontier technologies, likely because the firm's initial skill structure is more complementary to

[38]We leave out the skill composition as this is highly correlated with the task composition. In fact, when using the skill group shares rather than the task shares, we find that firms with a higher initial share of specialists tend to grow faster. Hence, we could look at skill shifts rather than task shifts and find a complementary story of higher skills being beneficial for a firm's ability to grow when adopting frontier technologies.

[39]Note that firms were asked a random subset of these three covariates and we impute the randomly missing items as described in Appendix F.1-F.3. Using the original items and a missing indicator for randomly missing items provides comparable results.

Table 5: Employment growth of frontier adopters between 2011 and 2016

Dependent variable: Change in firm-level employment $(\Delta \log N_f)$							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Initial occupational employment structure (reference: RC/RM jobs):							
NRC share	0.265***	0.229**	0.234**	0.274**	0.348***	0.417***	0.357***
	(0.078)	(0.107)	(0.108)	(0.114)	(0.119)	(0.130)	(0.136)
NRM share	-0.071	-0.034	0.027	0.095	0.112	0.189	0.228
	(0.161)	(0.158)	(0.140)	(0.153)	(0.165)	(0.147)	(0.163)
Adj R^2	0.067	0.122	0.132	0.207	0.242	0.247	0.271
Firm characteristics		✓	✓	✓	✓	✓	✓
Workforce structure			✓	✓	✓	✓	✓
Capital structure				✓	✓	✓	✓
Perceptions					✓		
Skill demands						✓	✓
Training							✓

Notes: $N = 172$. The average change in the dependent variable among the 172 frontier adopters is $\Delta \log N_f = 0.118$. The set of firm characteristics include a dummy for the self-assessment of whether the firm should be considered as a manufacturer or service provider, sector (5 categories), firm size (4 categories), firm age (dummy if founded after 2001), firm performance (AKM fixed effects). The workforce characteristics refer to the initial values in 2011 and include age shares (3 categories) and the share of IT workers. The capital structure characteristics include the initial level of technology-specific capital stocks (in logs) and the change in the technology-specific capital stocks (in logs) between 2011 and 2016. Perceptions include the assessment of chances and risks related to the adoption of frontier technologies (16 dimensions), see Appendix F.1. Skill demands include the assessment of how skill demands changed in the past five years (14 dimensions), see Appendix F.2. Training includes a set of variables describing how the extent and content of vocational and further training changed in the past five years (10 dimensions), see Appendix F.3.

new frontier technologies from the start compared to a firm with initially low NRC employment share. The downward bias in column 4 also tentatively suggests that firms with an initially high share of NRC workers select into more complex frontier technologies with more upskilling and training needs which diminish firms' growth potential.

This exercise demonstrates that the composition effect is robust to adding a wide set of controls. Firms who are already specialized in complementary, non-routine cognitive jobs, grow faster. Firms' NRC employment specialization therefore seems to play a crucial role in firms' ability to benefit from the adoption of new frontier technologies.

6 Conclusion

Extensive literature shows that computer-controlled machines of the third industrial revolution have replaced workers in routine jobs (Acemoglu and Autor, 2011). Yet, only limited evidence exists on whether and how technologies of the most recent wave of technological advances contribute to aggregate de-routinization, or, whether the adoption of frontier technologies results in different occupational employment shifts. We contribute to this debate by looking at the extent to which aggregate de-routinization is driven by technology adoption at the firm level. In particular, we decompose the aggregate de-routinization in Germany during the period 2011-2016 into changes arising within firms or due to employment reallocation between firms within versus between adopter groups. For the analyses, we created a novel linked employer-employee data that captures firms' actual technology investments for a representative sample of service and manufacturing firms. As one major advantage, the data distinguishes most recent frontier technologies from mature digital technologies that have been established since the 1970s.

Our study reveals several novel findings: First, we show that the adoption of frontier technologies tends to be routine-replacing in the economy, i.e., the aggregate routine share of the workforce employed at frontier adopters declines significantly and even more than for digital adopters and non-adopters.

Second, decomposing the aggregate decline in routine employment reveals a further surprising finding: Frontier technologies do not coincide with the widespread replacement of routine jobs within firms. In fact, the average frontier adopter (insignificantly) increases rather than decreases its routine employment share. Instead, the aggregate decline in routine jobs among the workforce in frontier firms stems from the heterogeneity between frontier firms: Initially larger frontier adopters reduce their routine employment share, i.e., increase their non-routine cognitive employment share, relatively more compared to smaller frontier adopters (scale effect). At the same time, initially less routine-intensive firms, i.e., more non-routine cognitive intensive firms, experience stronger employment growth (composition effect).

Finally, we delve deeper into the drivers of this scale and composition effect using firm-level regressions of occupation shifts and employment growth. The size-related heterogeneity in the routine employment decline (scale effect) most likely is explained by differences in how firms adopt frontier technologies which is reflected in changing skill demands and training needs. The growth-enhancing effect of an initially less routine-intensive firm workforce (composition effect) speaks for a strong role of firms' initial occupation specialization for its ability to benefit from the adoption of new frontier technologies. These findings indicate that the routine-replacing effect of frontier technologies hinges on having the right skills: either from the start or by complementary investments in upskilling.

Our paper comes with three implications: First, the diffusion of frontier technologies is likely to accelerate de-routinization as frontier adopters contribute the most to aggregate de-routinization. Second, the decline of routine-intensive employment is concentrated among a subset of frontier adopters, as opposed to widespread de-routinization within frontier firms across the economy. Our results thus suggest a novel dimension of a rising between-firm inequality among frontier firms. This so far unknown heterogeneity among frontier firms might be one reason why previous micro-level evidence did not find clear trends regarding the routine-replacing nature of these technologies. In fact, studies focusing on the within-firm adjustments after technology adoption likely underestimate the routine-replacing character of frontier technologies. Third, our results tentatively imply a potential self-reinforcing character of de-routinization: If a less routine-intensive workforce fosters firm growth and larger firms are more effective in replacing routine jobs, which then again fosters firm growth, this might result in continued and accelerating de-routinization on the one hand, and increasing market concentration on the other hand, especially in an early adoption period where complementary skills are scarce. As frontier technology becomes more mature, such non-routine skills may also get more widespread. According to our findings, this should counteract a rise in between-firm inequality, but further accelerate de-routinization. In fact, the de-routinization we observe results from declining shares of routine occupations, and thus, it should be considered as a lower bound since it does not capture task changes within occupations over time. In order to better understand such dynamics, future research should re-examine aggregate employment shifts as frontier technologies mature and take a look at the role of skill availability as a mediating factor.

References

Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, 67(2):251–333.

Acemoglu, D., Anderson, G. W., Beede, D. N., Buffington, C., Childress, E. E., Dinlersoz, E., Foster, L. S., Goldschlag, N., Haltiwanger, J. C., Kroff, Z., Restrepo, P., and Zolas, N. (2022a). Automation and the workforce: A firm-level view from the 2019 annual business survey. NBER Working Paper No. 30659, National Bureau of Economic Research.

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, 4:1043–1171.

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

Acemoglu, D., Lelarge, C., and Restrepo, P. (2020). Competing with robots: Firm-level evidence from France. *AEA Papers and Proceedings*, 110:383–88.

Acemoglu, D. and Restrepo, P. (2018). Modeling automation. *AEA Papers and Proceedings*, 108:48–53.

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

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2020). What are the labor and product market effects of automation? New evidence from France. CEPR Discussion Paper No. 14443, Centre for Economic Policy Research.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Reenen, J. V. (2020). The fall of the labor share and the rise of superstar firms. *Quarterly Journal of Economics*, 135(2):645–709.

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

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):1553–97.

Babina, T., Fedyk, A., He, A., and Hodson, J. (2022). Artificial intelligence, firm growth, and product innovation. Available at SSRN: http://dx.doi.org/10.2139/ssrn.3651052.

Bartel, Ichniowski, and Shaw (2007). How does information technology affect productivity? Plant-level comparisons of product innovation, process improvement, and worker skills. *Quarterly Journal of Economics*, 122(4):1721–1758.

Barth, E., Bryson, A., Davis, J. C., and Freeman, R. (2016). It's where you work: Increases in the dispersion of earnings across establishments and individuals in the United States. *Journal of Labor Economics*, 34(S2):67–97.

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.

Bechmann, S., Tschersich, N., Ellguth, P., Kohaut, S., and Florian, C. (2021). Technical report on the IAB establishment panel - Wave 28. Fdz-methodenreport, 07/2021 en, Institute for Employment Research.

Berlemann, M. and Wesselhöft, J.-F. (2014). Estimating aggregate capital stocks using the perpetual inventory method. *Review of Economics*, 65(1):1–34.

Bessen, J., Goos, M., Salomons, A., and van den Berge, W. (2020). Firm-level automation: Evidence from the Netherlands. *AEA Papers and Proceedings*, 110:389–393.

Bettiol, M., Capestro, M., Di Maria, E., and Furlan, A. (2019). Impacts of industry 4.0 investments on firm performance: Evidence from italy. "Marco Fanno" Working Paper No. 0233, Dipartimento di Scienze Economiche "Marco Fanno".

Böckerman, P., Laaksonen, S., and Vainiomäki, J. (2019). Does ICT usage erode routine occupations at the firm level? *LABOUR*, 33(1):26–47.

Bonfiglioli, A., Crinò, R., Fadinger, H., and Gancia, G. (2020). Robot imports and firm-level outcomes. CESifo Working Paper No. 8741, Center for Economic Studies and Ifo Institute.

Bresnahan, T. F., Brynjolfsson, E., and Hitt, L. M. (2002). Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence. *Quarterly Journal of Economics*, 117(1):339–376.

Brynjolfsson and Hitt (2003). Computing productivity: Firm-level evidence. *The Review of Economics and Statistics*, 85(4):793–808.

Brynjolfsson, E. and Hitt, L. M. (2000). Beyond computation: Information technology, organizational transformation and business performance. *Journal of Economic Perspectives*, 14(4):23–48.

Brynjolfsson, E. and Milgrom, P. (2013). Complementarity in organizations. In Gibbons, R. and Roberts, J., editors, *The Handbook for Organization Economics*, pages 11–55. Princeton University Press.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108:43–47.

Card, D., Heining, J., and Kline, P. (2013). Workplace heterogeneity and the rise of the West German wage inequality. *Quarterly Journal of Economics*, 128(3):967–1015.

Cathles, A., Nayyar, G., and Rückert, D. (2020). Digital technologies and firm performance: Evidence from Europe. EIB Working Papers No. 2020/06, European Investment Bank.

Cirillo, V., Fanti, L., Mina, A., and Ricci, A. (2022). New digital technologies and firm performance in the Italian economy. *Industry and Innovation*, 30(1):1–30.

Cortes, G. M. (2016). Where have the middle-wage workers gone? A study of polarization using panel data. *Journal of Labor Economics*, 34(1):63–105.

Cortes, G. M., Lerche, A., Schönberg, U., and Tschopp, J. (2023). Technological change, firm heterogeneity and wage inequality. IZA Discussion Paper No. 16070, Institute of Labor Economics.

Cortes, G. M. and Salvatori, A. (2019). Delving into the demand side: Changes in workplace specialization and job polarization. *Labour Economics*, 57:164–176.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6):3104–3153.

Dengler, K., Matthes, B., and Paulus, W. (2014). Occupational Tasks in the German Labour Market – An Alternative Measurement on the Basis of an Expert Database. Fdz-methodenreport, 12/2014, Institute for Employment Research.

DeStefano, T., Kneller, R., and Timmis, J. (2018). Broadband infrastructure, ICT use and firm performance: Evidence for UK firms. *Journal of Economic Behavior & Organization*, 155:110–139.

Dixon, J., Hong, B., and Wu, L. (2021). The robot revolution: Managerial and employment consequences for firms. *Management Science*, 67(9):5586–5605.

Dustmann, C., Ludsteck, J., and Schönberg, U. (2009). Revisiting the German wage structure. *Quarterly Journal of Economics*, 124(2):843–881.

Falck, O., Gold, R., and Heblich, S. (2014). E-lections: Voting behavior and the internet. *American Economic Review*, 104(7):2238–2265.

Falck, O., Heimisch-Roecker, A., and Wiederhold, S. (2021). Returns to ICT skills. *Research Policy*, 50(7):104064.

Freeman, R. B., Ganguli, I., and Handel, M. J. (2020). Within-occupation changes dominate changes in what workers do. *AEA Papers and Proceedings*, pages 394–399.

Gaggl, P. and Wright, G. C. (2017). A short-run view of what computers do: Evidence from a UK tax incentive. *AEJ: Applied Economics*, 9(3):262–94.

Ganzer, A., Schmidtlein, L., Stegmaier, J., and Wolter, S. (2021). Establishment history panel 1975-2019, Revised version v2 from April 2021. Fdz-datenreport, 16/2020 en, Institute for Employment Research.

Gathmann, C. and Grimm, F. (2023). Advances in digital technologies in Europe and their local labor market effects. Mimeo.

Genz, S., Gregory, T., Janser, M., Lehmer, F., and Matthes, B. (2021). How do workers adjust when firms adopt new technologies? IZA Discussion Paper No. 14626, Institute of Labor Economics.

Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in Europe. *AER: Papers and Proceedings*, 99(2):58–63.

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. *The Review of Economics and Statistics*, 100(5):753–768.

Harrigan, J. and Reshef, A. (2015). Skill-biased heterogeneous firms, trade liberalization and the skill premium. *Canadian Journal of Economics*, 48(3):1024–1066.

Harrigan, J., Reshef, A., and Toubal, F. (2021). The march of the techies: Job polarization within and between firms. *Research Policy*, 50(7):104008.

Humlum, A. (2019). Robot adoption and labor market dynamics. Working Paper, Available at: https://economics.yale.edu/sites/default/files/humlumjmp_111419.pdf.

Janser, M., Lehmer, F., and Zierahn-Weilage, U. (2023). Computing capital stocks in the German social security records and quantifying their role for wage inequality. Mimeo.

Kerr, S., Maczulskij, T., and Maliranta, M. (2020). Within and between firm trends in job polarization: The roles of globalization and technology. *Journal of Economic Geography*, 20(4):1003–1039.

Koch, M., Manuylov, I., and Smolka, M. (2021). Robots and firms. *The Economic Journal*, 131(638):2553–2584.

McElheran, K., Li, J. F., Brynjolfsson, E., Kroff, Z., Dinlersoz, E., Foster, L. S., and Zolas, N. (2022). AI adoption in America: Who, what, and where. NBER Working Paper No. 31788, National Bureau of Economic Research.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725.

Meyer, J. (2011). Workforce age and technology adoption in small and medium-sized service firms. *Small Business Economics*, 37:305–324.

Nicoletti, G., von Rueden, C., and Andrews, D. (2020). Digital technology diffusion: A matter of capabilities, incentives or both? *European Economic Review*, 128:103513.

Rammer, C., Fernández, G. P., and Czarnitzki, D. (2022). Artificial intelligence and industrial innovation: Evidence from German firm-level data. *Research Policy*, 51(7):104555.

Reijnders, L. S. M. and de Vries, G. J. (2018). Technology, offshoring and the rise of non-routine jobs. *Journal of Development Economics*, 135:412–432.

Seamans, R. and Raj, M. (2019). Artificial intelligence, labor, productivity, and the need for firm-level data. In *The Economics of Artificial Intelligence: An Agenda*, pages 553–565. University of Chicago Press.

Song, J., Price, D. J., Guvenen, F., Bloom, N., and Von Wachter, T. (2019). Firming up inequality. *Quarterly Journal of Economics*, 134(1):1–50.

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.

Webb, M. (2020). The impact of artificial intelligence on the labor market. Working Paper, Available at: https://www.michaelwebb.co/webb_ai.pdf.

Zolas, N., Kroff, Z., Brynjolfsson, E., McElheran, K., Beede, D. N., Buffington, C., Goldschlag, N., Foster, L., and Dinlersoz, E. (2020). Advanced technologies adoption and use by US firms: Evidence from the annual business survey. NBER Working Paper No. 28290, National Bureau of Economic Research.

Appendix

Appendix A Additional tables

Table A.1: Average technology shares as reported in the firm survey

	All	non-adopters	digital adopters	frontier adopters
	(1)	(2)	(3)	(4)
Panel A: Share of tools in 2011				
manual tech	0.573	0.788	0.409	0.371
digital tech	0.376	0.196	0.545	0.348
frontier tech	0.051	0.017	0.046	0.281
Panel B: Share of tools in 2016				
manual tech	0.500	0.776	0.284	0.262
digital tech	0.427	0.203	0.650	0.324
frontier tech	0.073	0.020	0.066	0.413

Notes: Table shows the subjective assessment of the survey respondent how the firm's tools are distributed over the technology classes as defined in Table 1. The shares are weighted with firm stratification weights. Panel A presents the results for 2011. Panel B presents the results for 2016.

Table A.2: Technology use across sectors

	Manufacturing		Service	
	(1)	(2)	(3)	(4)
Panel A: Technology-specific capital stock in 2011 (in 1,000 € and in %)				
manual tech	117.3	45.3%	63.8	47.4%
digital tech	125.1	48.3%	61.2	45.5%
frontier tech	16.5	6.4%	9.5	7.1%
Panel B: Technology-specific capital stock in 2016 (in 1,000 € and in %)				
manual tech	174.6	38.2%	99.5	38.4%
digital tech	231.4	50.7%	132.6	51.3%
frontier tech	50.7	11.1%	26.7	10.3%
Panel C: Representing number of firms and workers				
# firms (in K)	312.7		1612.4	
# workers (in M)	6.4		21.5	
Observations	583		1,077	

Notes: Table shows mean statistics weighted with firm stratification weights. Uneven columns give the capital stock in 1,000 € and even columns give the relative shares of the capital stocks across technology classes in % adding up to 100. Panel A presents the results for 2011. Panel B presents the results for 2016.

Table A.3: Initial firm and workforce characteristics by adoption group

	Non-adopters	digital adopters	frontier adopters
	(1)	(2)	(3)
Panel A: Firm characteristics			
Sector			
Secondary not know-int	0.456	0.160	0.090
Secondary know-int	0.017	0.025	0.008
Tertiary not know-int	0.404	0.481	0.363
Tertiary know-int	0.114	0.305	0.492
ICT sector	0.009	0.028	0.047
Manufacturing firm	0.244	0.108	0.038
Avg. firm employment (in heads)	10.5	17.4	19.1
Avg. firm daily wage (in €)	80.3	98.1	105.4
Firm Wage Premium	12.2	14.1	15.1
MDF port within reach	0.960	0.959	0.997
Panel B: Workforce characteristics			
Worker shares by main task of occupation			
NRC	0.238	0.329	0.358
RC/RM	0.399	0.449	0.491
NRM	0.362	0.222	0.151
Worker shares by required skill level			
Unskilled workers	0.139	0.100	0.119
Professionals	0.698	0.631	0.575
Specialists	0.088	0.130	0.136
Experts	0.075	0.139	0.170
Worker shares by age group			
<40 years of age	0.430	0.390	0.420
40-55 years of age	0.458	0.487	0.480
>55 years of age	0.112	0.123	0.100
Observations	525	963	172

Notes: The tables displays averages among firms across adoption groups for the initial observation period 2011 based on the entire sample of 1,660 firms. Statistics in Panel A are weighted with firm stratification weights. Appendix Table B.1 defines the five sector aggregates. The variable manufacturing firm stems from a self-assessment during the firm survey of whether the firm should be considered as a manufacturer or service provider. The firm wage premium corresponds to firm fixed effects for the period 2003-2010 based on the method by Abowd et al. (1999). These AKM FE are only available for 1,536 firms. Panel B represents the share of workers employed in each category and statistics are weighted by employment-weighted firm stratification weights. NRC: non-routine cognitive jobs; RC/RM: Routine cognitive/manual jobs; NRM: non-routine manual jobs.

Table A.4: Average marginal effects from a multiomial probit model

	Adopting digital technology	Adopting frontier technology
	(1)	(2)
Manufacturing firm	-0.152**	-0.103**
	(0.062)	(0.041)
Firm wage premium	0.003	-0.003
	(0.005)	(0.002)
Firm employment	0.026	0.021***
	(0.019)	(0.008)
MDF port within reach	-0.166	0.124**
	(0.102)	(0.052)
Worker shares by main task of occupation (reference: NRM)		
NRC	0.138	0.094**
	(0.089)	(0.047)
RC/RM	0.221***	0.065
	(0.070)	(0.042)
Worker shares by required skill level (reference: Professionals)		
Experts	0.334***	0.023
	(0.129)	(0.056)
Specialists	0.147	0.071
	(0.105)	(0.052)
Unskilled workers	-0.294**	0.071
	(0.145)	(0.065)
Worker shares by age group (reference: 40-55 years of age)		
<40 years of age	-0.024	-0.077*
	(0.091)	(0.041)
>55 years of age	0.253	-0.153
	(0.156)	(0.103)
Observations	1,660	1,660
Mean dep. variable	0.480	0.079

Notes: Columns show average marginal effects from a mutinomial probit in column (1) for being a 3.0-adopter compared to being a non-adopter and column (2) for being a 4.0 adopter compared to a non-adopter. The variable set is equivalent to the one used in column (3) of Table 3. Model estimated with firm stratification weights. NRC: non-routine cognitive jobs; RC/RM: Routine cognitive/manual jobs; NRM: non-routine manual jobs. Robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1.

Appendix B Data

Appendix B.1 Sector definition

To group firms along their business activities into sector aggregates, we follow the basic division along the primary (agriculture, fishing, or mining), secondary (manufacturing or construction), and tertiary sectors (trade, banking, education, or health services). Given that the primary sector only employs a minor share of total employment (about 1.5%), we focus on the distinction between the secondary and tertiary sectors. Instead, we further distinguish between knowledge-intensive and not knowledge-intensive sectors within the secondary and tertiary sectors and identify the ICT sector (information technology services or telecommunication services) separately. Firms in the primary sector are assigned to the not knowledge-intensive manufacturing sector.

Table B.1: Definition of sectors

Sector aggregates	German classification of economic activities
Secondary Sector	
Not knowledge-intensive manufacturing	11-99, 101-182, 221-259, 310-439
Knowledge-intensive manufacturing	191-212, 265-267, 271-309
Tertiary Sector	
Not knowledge-intensive service	451-464, 466-563, 681-683, 771-856, 871-889, 920-949, 952-990
Knowledge-intensive service	581, 591-602, 639-663, 691-750, 861-869, 900-910
ICT Sector	261-264, 268, 465, 582, 611-631, 951

Notes: The corresponding time-consistent three-digit codes of the German Classification of Economic Activities 2008 (WZ 2008) correspond to the NACE Rev. 2 industry codes.

Appendix B.2 Capital stock imputation

Although the IAB-ZEW-Labor Market 4.0-Establishment Survey (BIZA) includes rich information on firms' tools by level of technological sophistication, it misses information on the absolute value of the capital stock. Since the survey was conducted as a CATI, we considered it infeasible to ask for the technology-specific capital stocks across time during the interview. We adopt the procedure proposed by Janser et al. (2023) to impute capital stock information for the firms in our sample. The procedure consists of two steps.

In the first step, Janser et al. (2023) compute establishment-specific capital stocks using the investment data in the IAB Establishment Panel based on the perpetual inventory method. The IAB Establishment Panel is an annual representative survey on various topics including determinants of labor demand, investment activity, and workforce characteristics. The basis for the sampling of the IAB Establishment Panel comprises the universe of all German establishments with at least one employee subject to social insurance contributions as of the reference date 30 June of the previous year Bechmann et al. (2021). The procedure proposed by Janser et al. (2023) relies on the investment information from the IAB Establishment Panel for conducting a perpetual inventory approach following Berlemann and Wesselhöft (2014). The approach comprises several methods to remove outliers and control for business cycle fluctuations.

In the second step, Janser et al. (2023) merge these firms with additional administrative information that is available for the universe of firms and train a statistical learning model to predict capital stocks based on information that is available for all firms in Germany. The administrative information stems from the IAB Establishment History Panel. [40] The statistical learning method is divided into one model for firms' capital stock levels and one for changes in firms' capital stocks. The trained statistical learning models are used to predict capital stocks per year for all our survey establishments. As predictors for the capital stock K_{jt} for establishment j and year $t = 2011, ..., 2016$ several characteristics are used such as establishment size, age, sector affiliation, location, and employment structure.

It can be shown that those imputed capital stocks closely predict external capital stock information for a subset of firms for which direct information on firms' capital stocks is available. We use these predicted capital stocks per year for the survey establishments in our sample to assign each establishment a yearly overall capital stock K_{ft}. We compute technology-specific capital stocks K_{fgt} by multiplying the overall capital stock K_{ft} with the shares of firms' tools by technology level k_{fgt}.

Appendix B.3 Technology adoption definition

To classify firms according to their technology investments, we proceed in four steps: First, we classify a firm f as a g-adopter if it predominantly increases its g-technology capital stock relative to the other technology capital stocks:

$$g\text{-adopter} \Leftarrow \max\left(\Delta K_{f,g=1.0/2.0(non)}, \Delta K_{f,g=3.0(digital)}, \Delta K_{f,g=4.0(frontier)}\right) \tag{8}$$

where $g \in \{1.0/2.0, 3.0, 4.0\}$ indexes the level of technological sophistication and $\Delta K_{f,g}$ defines the change in the technology-specific capital stocks between 2011 and 2016 $K_{fg}^{t=2016} - K_{fg}^{t=2011}$.

[40] See Ganzer et al. (2021) for more details about the BHP. Access to the BHP is provided by the Research Data Center (FDZ) of the German Federal Employment Agency at the IAB (see http://fdz.iab.de).

For simplicity, we drop subscript f and write 2.0, whenever we mean 1.0/2.0. In particular:

$$g\text{-adopter}=\begin{cases} 4.0 & \text{if } \Delta K_{g=4.0} \geq \Delta K_{f,g=2.0} \quad \wedge \quad \Delta K_{g=4.0} \geq \Delta K_{g=3.0} \quad \wedge \quad \Delta K_{g=4.0} \neq 0 \quad \wedge \quad k_{g=4.0,t_{16}} \neq 0 \\ 3.0 & \text{if } \Delta K_{g=3.0} > \Delta K_{g=2.0} \quad \wedge \quad \Delta K_{g=3.0} > \Delta K_{g=4.0} \quad \wedge \quad \Delta K_{g=3.0} \neq 0 \quad \wedge \quad k_{g=3.0,t_{16}} \neq 0 \\ 2.0 & \text{if } \Delta K_{g=2.0} \geq \Delta K_{g=3.0} \quad \wedge \quad \Delta K_{g=2.0} > \Delta K_{g=4.0} \quad \wedge \quad \Delta K_{g=2.0} \neq 0 \quad \wedge \quad k_{g=2.0,t_{16}} \neq 0 \end{cases}$$

This step assigns all but two survey firms (after our selection steps) to one of the three adoption groups.

Second, if one of the other technology-specific capital shares is zero, the g-adoption status is defined as follows:

$$g\text{-adopter}=\begin{cases} 4.0 & \text{if } k_{g=4.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=3.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=2.0,t_{16}} = 0 \quad \wedge \quad \Delta K_{g=4.0} > \Delta K_{g=3.0} \\ 4.0 & \text{if } k_{g=4.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=2.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=3.0,t_{16}} = 0 \quad \wedge \quad \Delta K_{g=4.0} > \Delta K_{g=2.0} \\ 3.0 & \text{if } k_{g=3.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=2.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=4.0,t_{16}} = 0 \quad \wedge \quad \Delta K_{s=3.0} > \Delta K_{g=2.0} \\ 3.0 & \text{if } k_{g=3.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=4.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=2.0,t_{16}} = 0 \quad \wedge \quad \Delta K_{g=3.0} > \Delta K_{g=4.0} \\ 2.0 & \text{if } k_{g=2.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=3.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=4.0,t_{16}} = 0 \quad \wedge \quad \Delta K_{g=2.0} > \Delta K_{g=3.0} \\ 2.0 & \text{if } k_{g=2.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=4.0,t_{16}} \neq 0 \quad \wedge \quad k_{g=3.0,t_{16}} = 0 \quad \wedge \quad \Delta K_{g=2.0} > \Delta K_{g=4.0} \end{cases}$$

This step assigns an adoption status to the remaining 2 survey firms.

Difference to a capital-shares approach. Our approach defined above relies on the use of imputed technology-specific capital stocks K_{fgt}. Alternatively, one could rely on changes in the reported technology shares k_{fgt}. The former approach has the disadvantage of requiring imputed capital stocks, while its main advantage is that it accounts for the net investments and yields a more stringent classification of frontier adopters. The latter approach has the disadvantage of not taking into account the level of capital stocks and its main advantage is to not require additional assumptions related to the capital stock imputation. In the following, we describe the differences between the two approaches in detail.

Consider two firms A and B which raise their shares of frontier and digital technologies both by 1pp at the expense of older manual technologies. Suppose both firms initially rely on digital and frontier technologies to the same extent $k_{f,3.0,t=0} = k_{f,4.0,t=0}$. In this case, both approaches would assign firm A and B to the group of frontier adopters. Suppose the initial shares of digital technologies exceeded the share of frontier technologies for firm B $k_{B,3.0,t=0} > k_{B,4.0,t=0}$, while firm A still uses digital and frontier technologies to the same extent in the initial year $k_{A,3.0,t=0} = k_{A,4.0,t=0}$. In this case, the net investments into 3.0 technologies exceed the net investments into frontier technologies for firm B $I_{B,3.0} > I_{B,4.0}$, while they are equally large for firm A $I_{A,3.0} = I_{A,4.0}$. In this case, the capital stock approach assigns firm B as a digital adopter, while the technology share approach assigns firm B as a frontier adopter.

More formally, the relationship between both approaches can be expressed as follows:

$$\Delta K_{fg} = K_{f,t=0}\Delta k_{fg} + k_{fg,t=0}\Delta K_f + \Delta K_f \Delta k_{fg} \tag{9}$$

where $K_{f,t=0}$ is the total capital stock across all technology types, $k_{fg,t=0}$ is the g-specific technology share for firm f in the initial period $t = 0$. Δ refers to the difference between 2011 and 2016. If the initial technology shares were the same across all technology types ($k_{fg,t=0} = k_{fg',t=0}$), the technology type with the largest increase in the share would also be the technology type with the largest increase in the capital stock. Frontier technologies typically have a much smaller share initially (see Table 2). Thus, focusing on changes in technology shares Δk_{fg} results in a more lenient classification of frontier adopters: small changes in frontier technology shares are sufficient to classify a firm as a frontier adopter. Focusing on changes in technology-specific capital stocks ΔK_{fg} additionally emphasizes firms' net investments, yielding a more stringent classification of frontier adopters and thus differentiating potential differences between digital and frontier investments more precisely, avoiding a bias towards zero. In addition, de-investments in older manual technologies might appear as investments in frontier technologies by focusing exclusively on changes in the technology shares. Therefore, we focus on changes in technology-specific capital stocks ΔK_{fg} instead of changes in the technology shares Δk_{fg}.

In Appendix E, we use the approach focusing on the reported technology shares k_{fgt} to replicate the main results of our paper and report in Appendix Table E.1 a comparison of the assignment of firms into the three adoption groups between both approaches.

Appendix C Decomposition analysis: details

We extend the decomposition of Autor et al. (2020) and Acemoglu et al. (2020) by introducing sampling weights to account for the firm survey design, decomposing the residual changes in the covariance term and differentiating between more than two categories of firms.

Definitions

- λ_f – share of jobs in firm f with main task type $\lambda \in \{NRC, RM/RC, NRM\}$
- s_f – share of firm f in total employment, ignoring survey weights
- w_f – survey weight of firm f
- $F = \sum_{f \in F} w_f$ – number of firms (survey-weighted)
- $g \in (N, A^{3.0}, A^{4.0})$ – groups of firms
- $G = \sum_{f \in g} w_f$ – number of firms in group g (survey-weighted). Note: $N + A^{3.0} + A^{4.0} = F$
- $\tilde{s}_f = \frac{w_f s_f}{\sum_{f' \in F} w_{f'} s_{f'}}$ – employment share of firm f
- $\tilde{s}_g = \sum_{f \in g} \tilde{s}_f$ – share of group g in total employment
- $\bar{s} = \frac{1}{\sum_{f \in F} w_f} \sum_{f' \in F} \tilde{s}_{f'} = \frac{1}{F}$ – average employment share (= inverse number of firms)
- $\bar{s}_g = \frac{1}{\sum_{f \in g} w_f} \sum_{f' \in g} \tilde{s}_{f'} = \frac{\tilde{s}_g}{G}$ – average employment share of group g
- $\lambda = \sum_{f \in F} \frac{w_f s_f}{\sum_{f' \in F} w_{f'} s_{f'}} \lambda_f = \sum_{f \in F} \tilde{s}_f \lambda_f$ – aggregate employment share of jobs with main task type $\lambda \in \{NRC, RM/RC, NRM\}$
- $\bar{\lambda} = \frac{1}{F} \sum_{f \in F} w_f \lambda_f$ – firm-weighted average employment share of jobs with main task type $\lambda \in \{NRC, RM/RC, NRM\}$
- $\bar{\lambda}_g = \sum_{f \in g} w_f \lambda_f \frac{1}{F}$ – firm-weighted group-specific employment share of jobs with main task type $\lambda \in \{NRC, RM/RC, NRM\}$

We drop time indices for simplicity unless needed. Firm-weights w_f are time constant, and so are the numbers of firms.

Within firm variation and covariance

We decompose changes in aggregate employment shares of jobs with main task type $\lambda \in \{NRC, RM/RC, NRM\}$:

$$\begin{aligned}
\Delta\lambda =& \lambda_t - \lambda_{t_0} = \sum_{f\in F} \tilde{s}_{ft}\lambda_{ft} - \sum_{f\in F} \tilde{s}_{ft_0}\lambda_{ft_0} \\
=& \sum_{f\in F} \left(\tilde{s}_{ft}\lambda_{ft} - \tilde{s}_{ft_0}\lambda_{ft_0} + \Delta\lambda_f \frac{w_f}{F} - \Delta\lambda_f \frac{w_f}{F}\right) \\
=& \Delta\bar{\lambda} + \sum_{f\in F} \left(\tilde{s}_{ft}\lambda_{ft} - \tilde{s}_{ft_0}\lambda_{ft_0} - \lambda_{ft}\frac{w_f}{F} + \lambda f t_0 \frac{w_f}{F}\right) \\
=& \Delta\bar{\lambda} + \Delta \sum_{f\in F} \left(\lambda_f \tilde{s}_f - \lambda_f \frac{w_f}{F}\right) \\
=& \Delta\bar{\lambda} + \Delta \sum_{f\in F} \lambda_f \left(\tilde{s}_f - \frac{w_f}{F}\right) \\
=& \Delta\bar{\lambda} + \Delta \sum_{f\in F} \left(\lambda_f - \bar{\lambda}\right)\left(\tilde{s}_f - \frac{w_f}{F}\right) \qquad (10)
\end{aligned}$$

for the last step, note that $\sum_{f\in F} \bar{\lambda}(\tilde{s}_f - \frac{w_f}{F}) = 0$.

We can split up the within component by groups:

$$\Delta\bar{\lambda} = \sum_g \Delta\bar{\lambda}_g \qquad (11)$$

We can analogously split up the summation of the covariance part from $\sum_{f\in F}(\dots)$ into $\sum_g \sum_{f\in g}(\dots)$ to compute the covariance term separately for each group of firms.

Decomposing the covariance term

We decompose the covariance term of adopter group g:

$$\begin{aligned}
&\Delta \sum_{f\in g} \left(\lambda_f - \bar{\lambda}\right)\left(\tilde{s}_f - \frac{w_f}{F}\right) \\
&= \Delta \sum_{f\in g} \left[\left(\lambda_f - \bar{\lambda}\right)\left(\tilde{s}_f - \frac{w_f}{F}\right) - \left(\bar{\lambda}_g - \bar{\lambda}\right)\left(\tilde{s}_f - \frac{w_f}{F}\right) + \left(\bar{\lambda}_g - \bar{\lambda}\right)\left(\tilde{s}_f - \frac{w_f}{F}\right)\right] \\
&= \Delta \sum_{f\in g} \left[\left(\lambda_f - \bar{\lambda}_g\right)\left(\tilde{s}_f - \frac{w_f}{F}\right) + \left(\bar{\lambda}_g - \bar{\lambda}\right)\left(\tilde{s}_f - \frac{w_f}{F}\right)\right] \\
&= \Delta \sum_{f\in g} \left(\lambda_f - \bar{\lambda}_g\right)\left(\tilde{s}_f - \frac{w_f}{F}\right) + \Delta\left(\bar{\lambda}_g - \bar{\lambda}\right)\left(\tilde{s}_g - \frac{G}{F}\right) \qquad (12)
\end{aligned}$$

The first element is the within adoption groups effect, the second element is the between adoption groups effect.

Decomposing the within adoption groups effect

We further apply a shift-share decomposition to the within adoption groups effect:

$$\begin{aligned}
&\Delta \sum_{f\in g} \left(\lambda_f - \bar{\lambda}_g\right)\left(\tilde{s}_f - \frac{w_f}{F}\right) \\
=&\sum_{f\in g}\left(\tilde{s}_f - \frac{w_f}{F}\right)\Delta\left(\lambda_f - \bar{\lambda}_g\right) \\
&+\sum_{f\in g}\left(\lambda_f - \bar{\lambda}_g\right)\Delta\left(\tilde{s}_f - \frac{w_f}{F}\right) \\
&+\sum_{f\in g}\Delta\left(\lambda_f - \bar{\lambda}_g\right)\Delta\left(\tilde{s}_f - \frac{w_f}{F}\right)
\end{aligned} \tag{13}$$

For simplicity, we drop time indices – levels refer to initial levels, changes refer to changes between t and the initial time period. We denote the first element the (within adoption groups) scale effect, the second element the (within adoption groups) composition effect and the third element the (within adoption groups) simultaneous effect. We provide details on the interpretation of these three terms in the main paper.

Decomposing the between adoption groups effect

We analogously apply a shift-share decomposition to the between adoption groups effect:

$$\begin{aligned}
&\Delta\left(\bar{\lambda}_g - \bar{\lambda}\right)\left(\tilde{s}_g - \frac{G}{F}\right) \\
=&\left(\tilde{s}_g - \frac{G}{F}\right)\Delta\left(\bar{\lambda}_g - \bar{\lambda}\right) \\
&+\left(\bar{\lambda}_g - \bar{\lambda}\right)\Delta\tilde{s}_g \\
&+\Delta\left(\bar{\lambda}_g - \bar{\lambda}\right)\Delta\tilde{s}_g
\end{aligned} \tag{14}$$

Note that $\Delta\frac{G}{F} = 0$ by definition (constant firm weights). For simplicity, we drop time indices – levels refer to initial levels, changes refer to changes between t and the initial time period. We denote the first element the (between adoption groups) scale effect, the second element the (between adoption groups) composition effect, and the third element the (between adoption groups) simultaneous effect.

(1) We denote as '(between adoption groups) scale effect' employment shifts across the

three occupation groups λ that stem from differences in the initial average firm size between the three adoption groups in combination with differential average employment shifts across the three occupation groups λ among firms of the three adoption groups. As an example, if frontier adopters are larger than the average firm and at the same time decrease their routine job share, the initial size difference would amplify this effect and further decrease the overall routine employment share.

(2) We denote as '(between adoption groups) composition effect' employment shifts across the three occupation groups λ that result from differences in initial occupational-specific employment shares between the three adoption groups in combination with differential employment growth. For example, suppose frontier adopters have a higher routine employment share initially and at the same time experience lower employment growth than firms belonging to the other adoption groups. In that case, the '(between adoption groups) composition effect' would contribute to a declining overall routine employment share.

(3) We denote as '(between adoption groups) simultaneous effect' changes in firm size and changes in occupation-specific employment shares that take place simultaneously but differ on average between the adoption groups. If, as an example, frontier adopters grow faster and simultaneously reduce their routine employment share more pronounced than non-adopters and digital adopters, this reallocation between adoption groups is captured by the '(between adoption groups) simultaneous effect'.

Appendix D Decomposition analysis: additional results

Figure D.1: Sources of the between adoption groups effect for frontier technology adopters

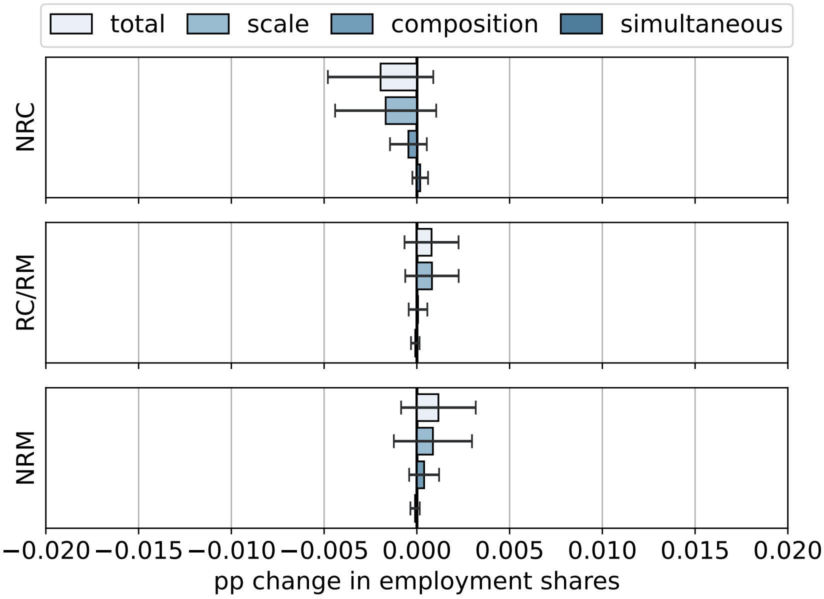

Notes: Decomposition of between-group component for frontier adopters as shown in equation (14): contribution of frontier adopters to aggregate de-routinization between 2011 and 2016 due to average differences between frontier adopters and other adoption groups related to size or composition. Scale, composition, and simultaneous components add up to the total between-group effect of frontier adopters. NRC: Non-routine cognitive tasks; RC/RM: Routine cognitive/manual tasks; NRM: Non-routine manual tasks. Confidence bands based on jackknife standard errors.

Figure D.2: Contribution of digital adopters to aggregate de-routinization

(a) Sources of de-routinization between 2011 and 2016 for digital adopters

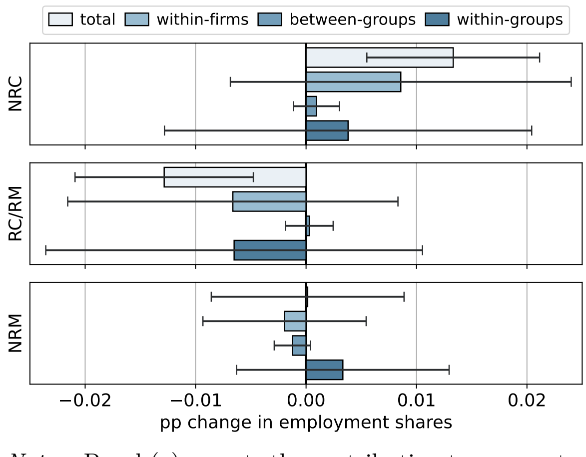

(b) Sources of the within adoption groups effect for digital adopters

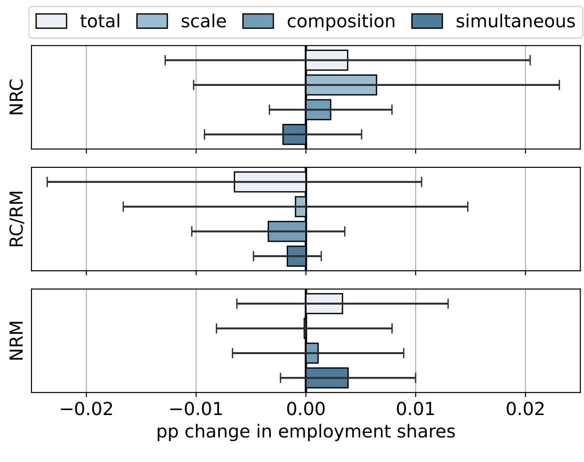

Notes: Panel (a) reports the contribution to aggregate de-routinization between 2011 and 2016 by digital adopters differentiating between the three components from equation 5. Panel (b) reports a decomposition of within adoption groups component for digital adopters as shown in equation (13): contribution of digital adopters to aggregate de-routinization between 2011 and 2016 due to shifts induced between digital adopting firms related to size, composition differences, and a combination of both. Scale, composition, and simultaneous components add up to the total within adoption groups effect of digital adopters shown in Panel (a). NRC: non-routine cognitive jobs; RC/RM: routine cognitive/manual jobs; NRM: non-routine manual jobs. Confidence bands based on jackknife standard errors.

Figure D.3: Contribution of non-adopters to aggregate de-routinization

(a) Sources of de-routinization between 2011 and 2016 for non-adopters

(b) Sources of the within adoption groups effect for non-adopters

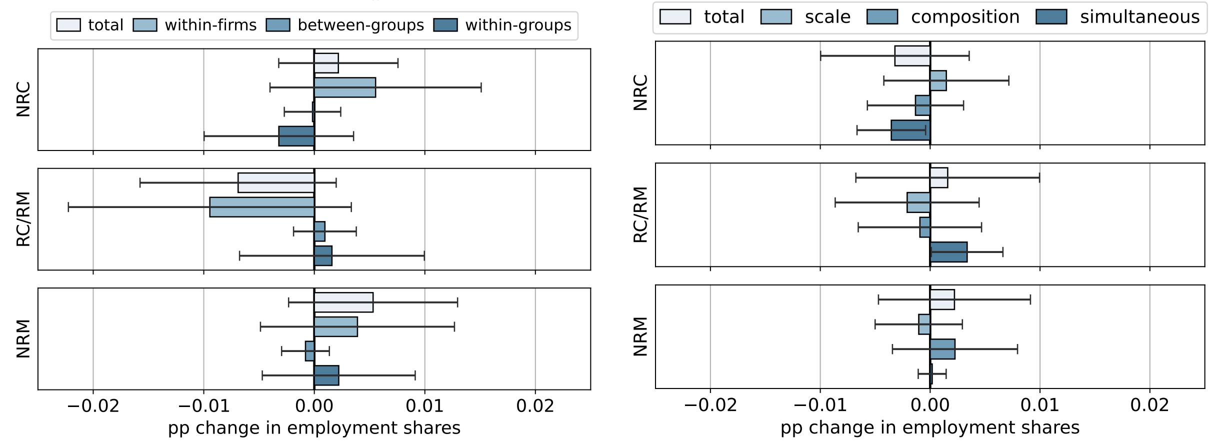

Notes: Panel (a) reports the contribution to aggregate de-routinization between 2011 and 2016 by non-adopters differentiating between the three components from equation 5. Panel (b) reports a decomposition of within adoption groups component for non-adopters as shown in equation (13): contribution of non-adopters to aggregate de-routinization between 2011 and 2016 due to shifts induced between non-adopting firms related to size, composition differences, and a combination of both. Scale, composition, and simultaneous components add up to the total within adoption groups effect of non-adopters shown in Panel (a). NRC: non-routine cognitive jobs; RC/RM: routine cognitive/manual jobs; NRM: non-routine manual jobs. Confidence bands based on jackknife standard errors.

Figure D.4: Skill shifts separately for each adoption group

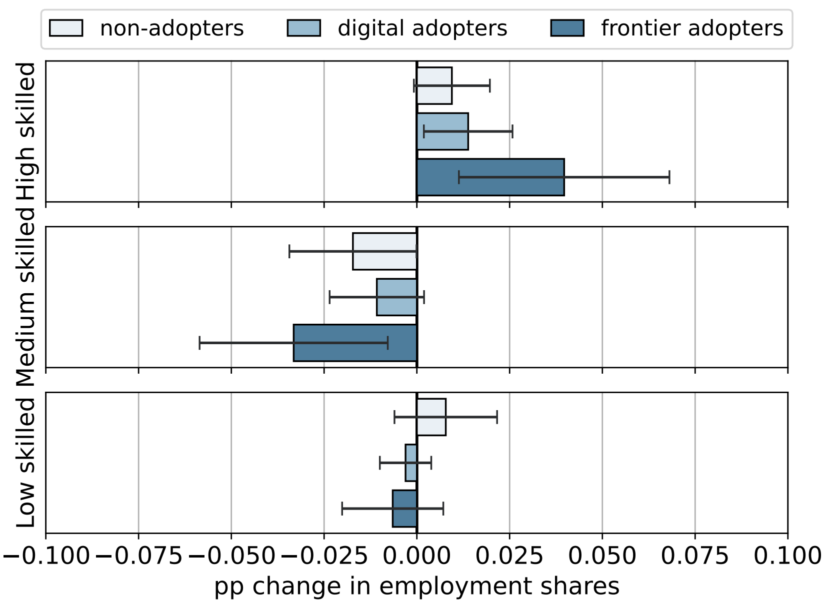

Notes: Changes in the skill-specific employment shares between 2011 and 2016 separately for each adoption group g. High-skilled refers to individuals with a university degree, medium-skilled refers to workers with a vocational training qualification, and low-skilled refers to individuals without a vocational training degree. Confidence bands based on jackknife standard errors.

Figure D.5: Employment shifts across requirement level groups separately for each adoption group

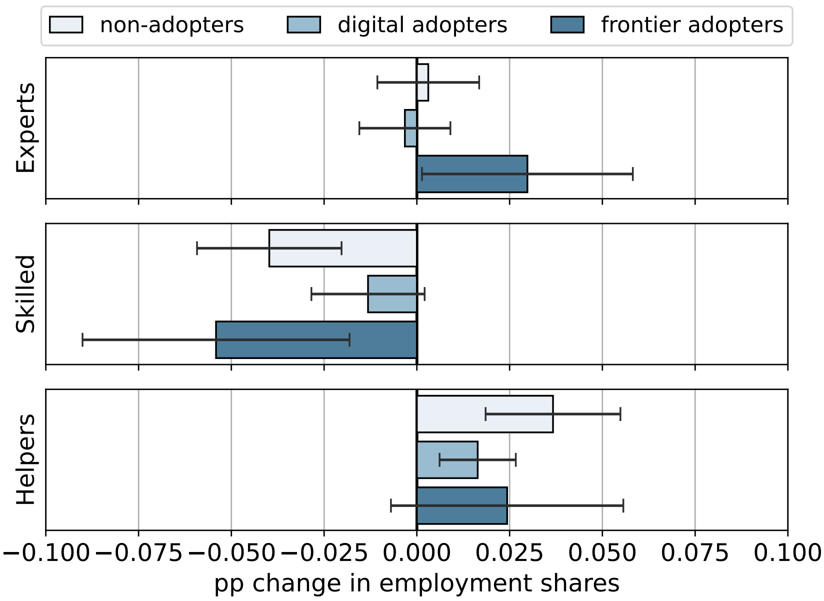

Notes: Changes in the employment shares across requirement level groups between 2011 and 2016 separately for each adoption group g. The three groups refer to different job requirement categories as they are formally distinguished in the German occupation classification (KldB-2010, 5-digit). Experts and specialists refer to the most complex work activities on their jobs. Skilled workers conduct somewhat complex activities while unskilled workers conduct simple and non-complex activities. Confidence bands based on jackknife standard errors.

Figure D.6: Employment shifts across wage groups separately for each adoption group

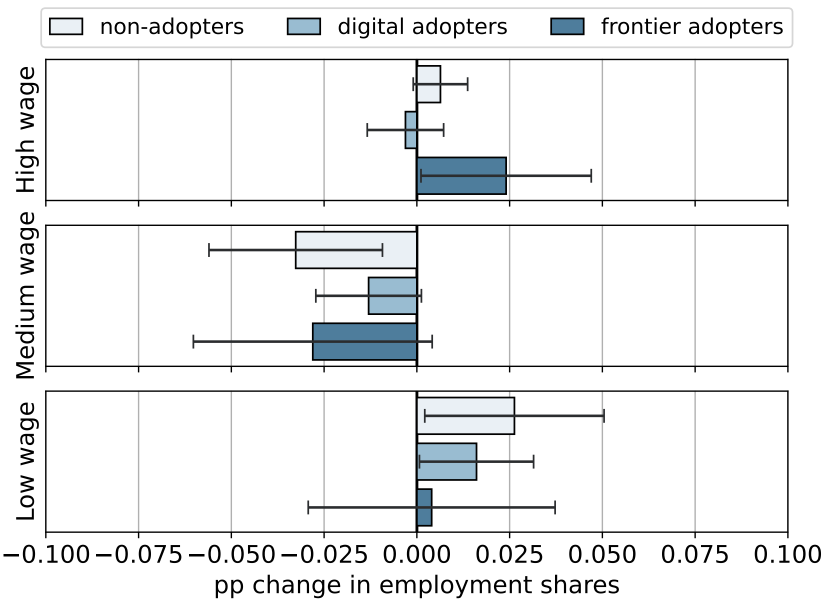

Notes: Changes in the employment shares across wage groups between 2011 and 2016 separately for each adoption group g. High wage occupations are based in the 75th percentile or higher of the median wage distribution of all occupations in Germany in the year 2012. Medium wage occupations are ranked between the 25th and 75th percentile of the median wage distribution of all occupations in Germany in the year 2012. Low wage occupations are ranked below the 25th percentile of the median wage distribution of all occupations in Germany in the year 2012. Confidence bands based on jackknife standard errors.

Appendix E Decomposition analysis: robustness

We run a number of robustness checks on the decomposition of the employment de-routinization in Germany between 2011 and 2016. In particular, we repeat the decomposition analysis using alternative sample choices and definitions diverging from those of the main sample as explained in section 2.2. We test whether the key findings are robust to (a) using a different time span that is unaffected by a structural break of the occupational classification in 2011, (b) the exclusion of non-profit firms, (c) the exclusion of outliers in employment growth or decline, (d) an alternative classification of firms into adoption groups as used in Genz et al. (2021), (e) an alternative classification of firms into adoption groups using only contemporaneous information on the technology use in 2016 and (f) an alternative task measure based on the O*NET data. For each robustness check, we show the results corresponding to Figure 2 including the within-firm effect and Figure 3 for comparison with the main sample:

(a) **Different time span.** We decompose the employment changes across the three occupational groups between 2012 and 2016 rather than between 2011 and 2016 as there was a change in the occupational classification between 2011 and 2012 that might result in a structural break. Resulting task changes are smaller because of the shorter time span, as expected, but the structural shifts are the same.

Figure E.1: De-routinization with new base year

(a) Contribution of frontier adopters (similar to Figure 2)

(b) Between firms within frontier adopter group (similar to Figure 3)

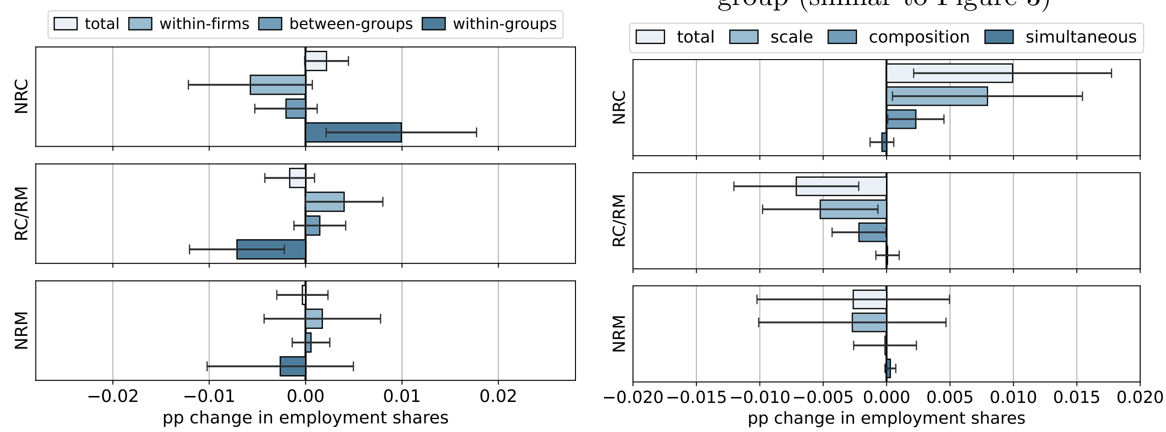

(b) **Excluding firms without revenues.** We exclude 110 firms that report to make no revenues. These establishments either belong to the public sector or are a non-profit organization and might differ from the private, profit-oriented sector. When excluding these firms, we still find similar result patterns as for the main sample:

Figure E.2: De-routinization excluding firms without revenues

(a) Contribution of frontier adopters (similar to Figure 2)

(b) Between firms within frontier adopter group (similar to Figure 3)

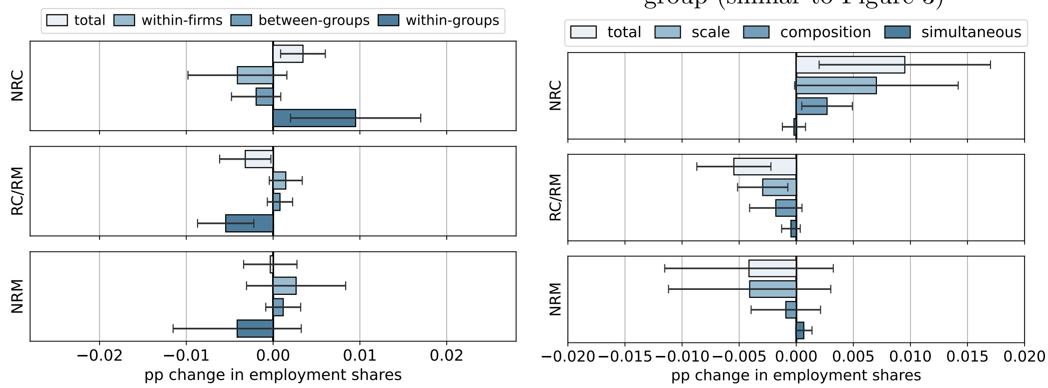

(c) **Excluding outliers in employment growth.** We exclude 24 establishments for which we observe outliers in employment growth across time. Such outliers may results from the in- or outsourcing of certain parts of the establishment which we cannot observe in the data. We thus exclude outliers, i.e., firms for which the annual employment growth in any of the years between 2011 and 2016 exceeds the 99th percentile plus one standard deviation (which corresponds to a growth factor of 4.4). We also exclude firms that lose more than 60% of their workforce (corresponding to the 1st percentile of the employment growth distribution) if the initial workforce size exceeds ten workers. Again, we find similar results.

Figure E.3: De-routinization excluding outliers in employment growth

(a) Contribution of frontier adopters (similar to Figure 2)

(b) Between firms within frontier adopter group (similar to Figure 3)

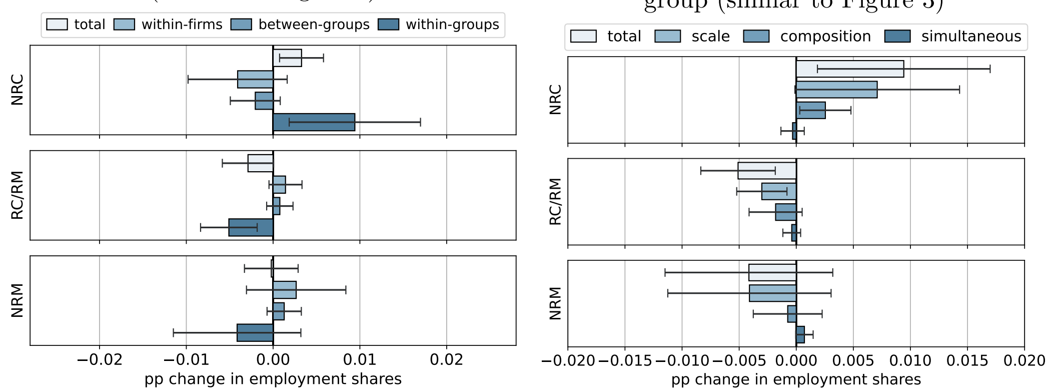

(d) **Assigning adoption groups based on technology shares.** Rather than using the changes in net investments ΔI_{fg} to classify establishments, we classify in this robustness check establishments into non-adopters, digital and frontier adopters based on percentage changes in the technology-specific capital shares k_{fgt} as reported from the survey and shown in Appendix Table A.1. In this alternative classification, an establishment would be considered as frontier adopter if the share of frontier technologies increased more than the shares of other technology types. Table E.1 compares this alternative adoption group assignment (based on changes in technology shares Δk_{fgt}) with assignment used in the main analysis (based on changes in net investments ΔI_{fg}). See also Appendix Section B.3 for a detailed comparison between both approaches.

Table E.1: Comparison between original and alternative adoption group assignments based on changes in technology shares

		Assignment based on technology shares (alternative)			
		non-adopters	digital adopters	frontier adopters	total
Assignment based on net investments (main analysis)	non-adopters	267	158	100	525
	digital adopters	7	712	244	963
	frontier adopters	0	7	165	172
	total	274	877	509	1,660

The comparison demonstrates that taking account of the capital stock and corresponding net investments results in a more restrictive assignment to the adopter groups. While there are only 172 frontier adopters in our main sample of analysis, using the capital shares as the basis for group assignment results in 509 adopters of frontier technologies. Our decomposition results nevertheless remain qualitatively robust.

Figure E.4: De-routinization when assigning adoption groups based on technology shares

(a) Contribution of frontier adopters (similar to Figure 2)

(b) Between firms within frontier adopter group (similar to Figure 3)

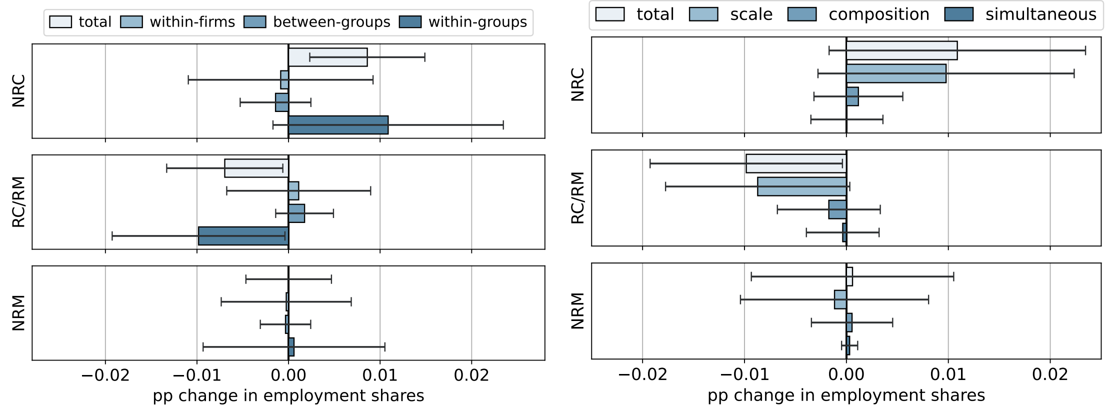

One minor difference relates to the source of the between firm task shifts occurring within the group of frontier adopters. With the more generous assignment of a frontier adoption status, the heterogeneity between frontier adopters is mostly driven scale effects while the composition effect contributes less to the heterogeneity among firms belonging to the frontier adopter group.

(e) **Assigning adoption groups based on 2016 technology use.** We now define our adoption groups based on the technology use in 2016 and refrain from using the retrospective information for the year 2011. For this, we define a firm to be a frontier adopter if the frontier technologies share in 2016 exceeds the 5th percentile of the corresponding distribution. Among the remaining firms, those whose digital technology share in 2016 exceeds the 5th percentile in the corresponding distribution are considered as digital adopters. All others are considered as non-adopters. We thus do not make any use of the retrospective information.

Table E.2: Comparison between original and alternative adoption group assignments based on technology use in 2016

		Assignment based on technology use in 2016 (alternative)			
		non-adopters	digital adopters	frontier adopters	total
Assignment based on net investments (main analysis)	non-adopters	417	91	17	525
	digital adopters	62	803	98	963
	frontier adopters	10	7	155	172
	total	489	901	270	1,660

Table E.2 shows that this alternative adoption group assignment is quite similar to the assignment based on net investments between 2011 and 2016. Moreover, decomposition results confirm the previous result pattern. Hence, results in our main analysis do not hinge on using the potentially unreliable, retrospective information.

Figure E.5: De-routinization when assigning adoption groups based on 2016 technology use

(a) Contribution of frontier adopters (similar to Figure 2

(b) Between firms within frontier adopter group (similar to Figure 3)

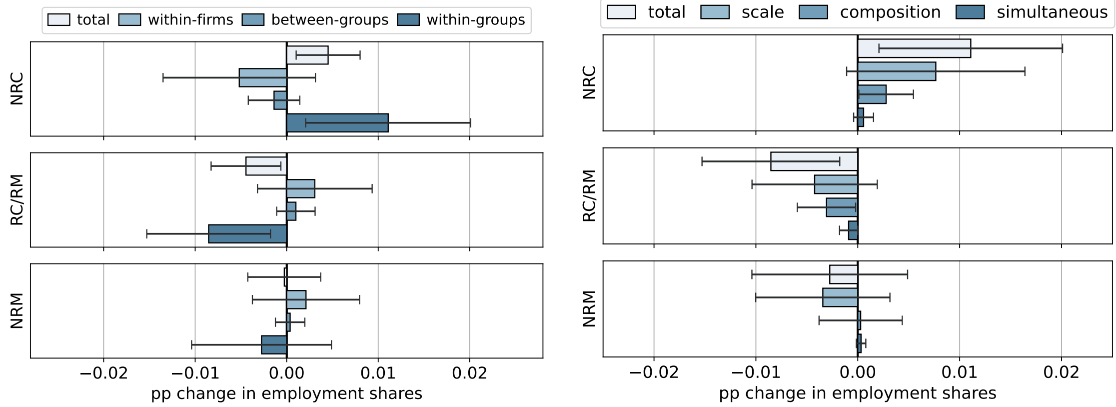

Notes: Panel (a) reports the contribution to aggregate employment shifts across task groups between 2011 and 2016 by frontier adopters differentiating between the three components from equation 5. Panel (b) reports a decomposition of within adoption groups component for frontier adopters as shown in equation (13): contribution of frontier adopters to aggregate employment shifts between 2011 and 2016 due to employment shifts induced between frontier adopting firms related to size, composition differences, and a combination of both. Scale, composition and simultaneous components add up to the total within adoption groups effect of frontier adopters shown in Panel (a). Confidence bands based on jackknife standard errors.

(f) **Main task assignment based on O*NET database.** Our classification of occupations is based on the main task at the 3-digit occupation level as reported from the German expert data base BERUFENET of the Federal Employment Agency (see section 2). As a robustness check, we classify firms according to the main task based on the task structure reported in the O*NET database. Since we have to conduct a crosswalk from the SOC occupation classification to the German KldB 2010, this alternative task measure introduces noise and accordingly leads to less strong task shifts, but the overall patterns remain.

Figure E.6: De-routinization with main task assignment based on O*NET database

(a) Contribution of frontier adopters (similar to Figure 2)

(b) Between firms within frontier adopter group (similar to Figure 3)

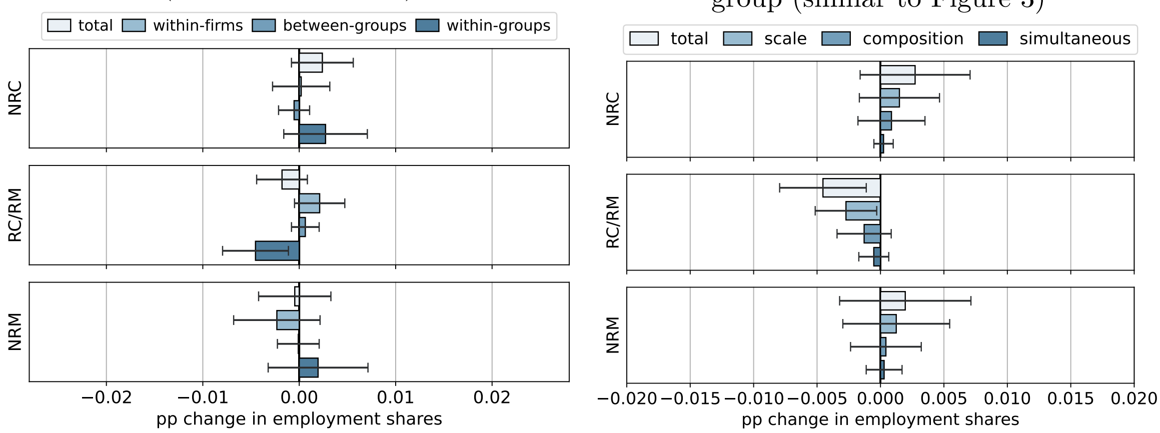

Appendix F Delving into the heterogeneity among frontier adopters

Appendix F.1 Perceived chances and risks of frontier technologies

The survey contains a set of questions on how firms perceive the opportunities and risks of using frontier technologies. The exact wording of the question is: “If you think about the opportunities and risks linked to these technologies in an establishment like yours. What do you think is true?”. Firms were selected randomly to answer 8 out of the following 16 questions, see Appendix Table F.1. On average, each firm answered 7.85 questions and 91% answered all 8 questions they were randomly assigned to. As a result, we have between 782 to 830 non-missing answers for each question. Each question was answered using a 5-point scale from (1) this is absolutely true, (2) this is more likely to be the case (3) neither, (4) this rather does not apply, (5) this does not apply at all.

To use the firm's perceptions as regressors in section 5, we impute the missing answers for each firm by exploiting the fact that each firm answered 8 out of 16 related questions that are likely to contain valuable correlated information. More precisely, we use a random forest classification for the binary information whether a firm fully agrees with the respective question or not.[41] We use the answers from all other perception measures as features. This is possible because we also include an additional missing dummy in case the respective perception measure is missing for firm f. Since the reason for missing perceptions is predominantly due to randomization, this is a feasible approach. In addition, the set of features includes firm and workforce characteristics and the firm's capital structure as in Table 4. In total, we use 120 features. We run the random forest algorithm on the full sample of all firms for which we observe the dependent variable irrespective of the firm's adoption status. Hence, the model also controls for adoption status to allow for a related heterogeneity in perceived chances and risks of adopting frontier technology. We further control for size to allow an additional size-related difference in perceptions.

Based on this model, we predict whether a firm fully agrees with each statement, and use this prediction for the regression analyses in section 5. Appendix Table F.1 includes several indicators for the observed perceptions and their imputed predictions including the out-of-bag (OOB) error as an indicator of the model's predictive power. Overall, Appendix Table F.1 suggests a good fit of our imputation model.

[41] Fine-tuning resulted in a random forest with 500 trees and several features per split which is equal to the square root of the total features, which resulted in the best out-of-sample error rate.

Table F.1: Questions on chances and risks related to frontier technologies

"The use of frontier technologies..."					**predicted prob for score=1**		
this is absolutely true (score=1) ... this does not apply at all (score=5)		**obs** (1)	**mean score** (2)	**share with score** = 1 (3)	**non-missing** N=obs (4)	**missing** N=1660-obs (5)	**OBB error** N=1,660 (6)
Chances							
labprod	... increases labor productivity	830	2.45	0.285	0.289	0.306	0.311
indprod	... enables customized products/services	814	2.61	0.273	0.284	0.308	0.305
newprod	... enables new products/services	813	2.77	0.223	0.235	0.279	0.257
transcostr	... reduces transport/warehousing costs	787	3.40	0.159	0.172	0.191	0.201
labcostr	... reduces labor costs	809	3.36	0.097	0.111	0.140	0.155
bodystrainr	... reduces workers's physical strain	782	3.47	0.081	0.094	0.126	0.132
energycr	... reduces energy costs	818	3.75	0.039	0.048	0.059	0.075
Risks							
cybercosti	... increases expenditures for cybersecurity	820	1.82	0.582	0.586	0.588	0.315
trainneed	... increases training needs	819	2.08	0.421	0.424	0.433	0.365
invcost	... comes with high investment costs	834	2.16	0.364	0.362	0.349	0.356
traincont	... changes the training content	830	2.46	0.347	0.355	0.374	0.342
extdep	... increases dependence on external firms	838	2.56	0.325	0.318	0.311	0.263
mentstraini	... increases worker's mental strain	784	2.98	0.225	0.215	0.189	0.156
restruc	... necessitates complex reorganization of firm	802	3.05	0.175	0.193	0.232	0.234
labshort	... is hampered by a shortage of skilled labor	849	3.15	0.171	0.184	0.197	0.187
econriski	... increases economic risks of firm	802	3.35	0.143	0.140	0.129	0.089

Notes: Perceptions are observed for the number of firms in column (1). Column (2) gives the observed average value of the 5-point scale for each perception, while column (3) reports the share of firms that fully agree (score=1). Columns (4)-(5) report the results from the predictions based on the random forest classification, i.e., the probability to fully agree (score=1) for the sub-sample of firms for which we also observe the perception (column 4) and the sub-sample of firms for which perceptions are missing (column 5). Column (6) displays the Out-Of-Bag error. All statistics are weighted with firm-stratification weights.

Appendix F.2 Skill demands

The survey contains a set of questions asking firms to assess whether they consider certain skill requirements to have decreased, remained stable, or increased. The exact wording of the question is: "If you think about the requirements that a job in your company entails. Please estimate which of the following requirements have gained or lost importance in the last five years." For time reasons, firms were randomly selected to answer this question for half of the 14 skill types listed in Appendix Table F.2. On average, each firm answered 6.98 and 98.5% answered the question for all 7 items they were randomly assigned to. For each item, firms could assess based on a 5-point scale whether the respective skill (1) strongly declined, (2) somewhat declined, (3) neither declined nor increased, (4) somewhat increased, and (5) strongly increased in importance during the last five years. Additionally, a firm could also answer that the skill is irrelevant in their firm.

To use these assessments as regressors in section 5, we impute the missing items for each firm by exploiting the fact that each firm answered about half of the 14 items which are likely to contain valuable correlated information. More precisely, we use a random forest to classify whether a firm considers the importance of the respective skill to have strongly increased or not.[42] For this, we use all other skill requirements as features of the model. This is possible because we also include an additional dummy in case the respective item is missing for firm f. Since the reason for missing items is predominantly due to randomization, this is a feasible approach. Moreover, we add estimated probabilities that certain skills are irrelevant.[43] In addition, the set of features includes firm and workforce characteristics as well as the firm's capital structure as in Table 4. We also add interactions between the adoption status and the initial employment shares across occupational groups. In total, we use 126 features in the random forest classification.

Based on this model, we predict the probability that a firm considers a certain skill requirement to have increased strongly in importance during the last five years. We use these predictions in section 5. The indicators for the observed and predicted skill requirements in Appendix Table F.2, as well as the Out-Of-Bag error rate suggest a good fit of our imputation model.

[42] Fine-tuning resulted in a random forest with 300 trees and several features per split which is equal to the square root of the total features, which resulted in the best out-of-sample error rate.

[43] For this, we first run a random forest classifier to estimate the probability that a certain skill is considered irrelevant by the firm. For this classification, we use the information on the irrelevance of the other observable skills (plus dummies if the answer is missing due to randomization of items). We do so only for a subset of skills for which at least 4% of all firms report the skill to be irrelevant in their firm. This only applies to *physstress*, *mandext*, *creativity*, *envknow*, *developIT*, and *interdisc*.

Table F.2: Questions on changes in skill demands

"The importance of [*skill*] during the last five years has ..."					**predicted prob for score=5**		
strongly declined (score=1) ... strongly increased (score=5)		**obs** (1)	**mean score** (2)	**share with score** = 5 (3)	**non-missing** N=obs (4)	**missing** N=1660-obs (5)	**OBB error** N=1,660 (6)
knowIT	knowledge in the application of IT	815	4.21	0.371	0.325	0.279	0.271
developIT	development of IT	822	4.07	0.255	0.173	0.144	0.233
learning	learning new skills and competencies	842	4.03	0.225	0.215	0.212	0.212
multitask	multitasking	842	3.83	0.209	0.198	0.209	0.199
dealcustom	dealing with customers	828	3.73	0.193	0.200	0.210	0.184
mentstress	working under high mental stress	829	3.74	0.183	0.174	0.189	0.167
envknow	environmental protection knowledge	830	3.85	0.182	0.148	0.122	0.163
indwork	independent working	800	3.71	0.151	0.171	0.179	0.186
creative	creativity	832	3.60	0.126	0.118	0.113	0.102
physstress	working under high physical stress	843	3.31	0.124	0.096	0.089	0.062
knowproc	process know-how	849	3.72	0.122	0.128	0.155	0.188
interdisc	interdisciplinary working	801	3.65	0.121	0.110	0.115	0.144
leadersh	leadership skills	836	3.62	0.120	0.127	0.133	0.148
mandext	manual dexterity	814	3.08	0.025	0.022	0.032	0.023

Notes: The change in importance of the skill requirements is observed for the number of firms in column (1). Column (2) gives the observed average value of the 5-point scale as described above, while column (3) reports the share of firms that suggest the respective skill requirement to have gained strongly in importance (score=5). Columns (4)-(5) report the results from the predictions based on the ordered logit model, i.e. the probability to consider the skill requirement to have gained strongly in importance (score=5) for the sub-sample of firms for which we also observe the perception (column 4) and the sub-sample of firms for which perceptions are missing (column 5). Column (6) displays the Out-Of-Bag error. All statistics are weighted with firm-stratification weights.

Appendix F.3 Training needs

The survey contains a set of questions asking firms to assess how vocational and further training changed in the five past years with respect to a number of different dimensions. In particular, firms are asked whether they agree to 10 different dimensions on a 5-point scale from (1) this is absolutely true, (2) this is more likely to be the case (3) neither, (4) this rather does not apply, (5) this does not apply at all. Appendix Table F.3 includes the statements firm's were asked to assess. However, firms were only asked to answer the questions if they offered vocational training and further training within the last five years. Accordingly, in contrast to firms' perceptions and skill requirements, items are not missing at random. However, using imputed training items vs. training items and a dummy indicator when those items are missing shows that our results in the main paper remain robust.

In order to use these assessments as regressors in section 5, we impute the missing items for each firm by using the information on all other training measures (including a dummy in case the respective item is missing for firm f). In addition, the set of features includes firm and workforce characteristics as well as the firm's capital structure as in Table 4. We run the random forest algorithm on the full sample of all firms for which we observe the dependent variable irrespective of the firm's adoption status. [44] Hence, the model also controls for adoption status to allow for a related heterogeneity in perceived development in vocational and further training. We further control for size to allow an additional size-related difference in perceptions. In total, we use 85 features in the random forest classification.

Based on this model, we predict the probability that a firm fully agrees with the respective statement about trends in training. We use these predictions in section 5. The indicators for the observed and predicted training indicators in Appendix Table F.3 suggests a good fit of our imputation model.

[44] Fine-tuning resulted in a random forest with 500 trees and several features per split which is equal to the square root of the total features, which resulted in the best out-of-sample error rate.

Table F.3: Questions on changes in vocational and further training needs

"Which of the following statements do you agree/disagree with?"					**predicted prob for score=1**		
this is absolutely true (score=1) ... this does not apply at all (score=5)		**obs** (1)	**mean score** (2)	**share with score = 1** (3)	**non-missing** N=obs (4)	**missing** N=1660-obs (5)	**OBB error** N=1,660 (6)
If you think of the ***vocational training*** *in your firm.*							
voctinterc	We increasingly train intercurricular skills.	1016	2.16	0.34	0.320	0.255	0.227
voctICT	We increasingly train ICT-skills.	1026	2.44	0.32	0.319	0.317	0.205
voctcont	The training content has changed in the last 5 years.	1003	2.54	0.25	0.235	0.164	0.179
voctmism	The mismatch between required and actual skills increases.	1010	3.34	0.14	0.136	0.124	0.059
voctnew	We train in other fields/occupations than in the past.	1027	4.03	0.04	0.051	0.057	0.077
If you think about ***further training*** *in your firm.*							
traincost	The costs for further training have increased.	1291	2.30	0.29	0.297	0.292	0.285
trainICT	We increasingly train ICT skills .	1295	2.65	0.26	0.253	0.220	0.205
trainhqual	Our workers can more easily get higher qualifications.	1296	3.08	0.15	0.146	0.116	0.160
trainelearn	We increasingly use E-Learning tools.	1303	3.38	0.12	0.137	0.121	0.170
traininterc	We increasingly train intercurricular skills.	1303	3.00	0.11	0.118	0.113	0.142

Notes: The assessment of the respective statement is observed for the number of firms in column (1). Column (2) gives the observed average value of the 5-point scale as described above, while column (3) reports the share of firms that fully agree with this statement (score=1). Columns (4)-(5) report the results from the predictions based on the random forest classification, i.e. the probability to fully agree to the respective statement (score=1) for the sub-sample of firms for which we also observe the answer (column 4) and the sub-sample of firms for which the answer is missing (column 5). Column (6) displays the Out-Of-Bag error. All statistics are weighted with firm-stratification weights.