Labor Market Effects of Automation: A Scoping Review

Anke Hassel, Didem Özkiziltan, Kilian Weil

September 16, 2022

Abstract

We review the literature of labor market effects of automation on employment and wages with a special focus on heterogeneous outcomes for skill groups. Limiting our scope to the past two OECD decades, the study includes contributions with at least one empirical specification explaining employment and wage outcomes due to technological change. By quantifying the body of literature according to (1) automation type, (2) empirical modelling strategies, and (3) heterogeneous effects for skill groups, we report conclusive evidence for high (middle) skilled segments with positive (negative) labor market outcomes, but find mixed evidence for low-skilled segments, and the economy as a whole. Next, we show that the variation of evidence is not random, but a function of research design choice. In particular, negative employment outcomes correlate with identification strategies that exclusively rely on occupational data. Lastly, we showcase best-practice designs and identify promising avenues for further research.

Contents

List of Figures

List of Tables

1 Introduction

We take stock of the research literature of the impact of technology on OECD labor markets by quantifying research findings of the last two decades. We limit our review to two labor market outcomes, and focus on 1) employment, and 2) wage effects of automation. Quality research should be generalizable, thus we provide benchmark criteria, which help to differentiate robust evidence from future scenarios (i.e. *risk* of automation research). Within the form of a *scoping review*, we distinguish different forms of technology (AI, ICT, Robots) and collect existing research findings according to their impact on different labor market skill brackets (the bias hypotheses). We add to the literature in two ways. First, by highlighting that different identification strategies are associated with different results. Second, by pointing to a gap in the literature which neglects the role of labour market institutions.

The paper is motivated by the far-reaching implications that are often associated with technologies for labour market outcomes. The topic has gained widespread attention across academic disciplines, from the public, as well as high-level policy discussions in OECD countries. Governments are called to design policies that can both exploit the potential of new technologies for growth and innovation and make sure that employees benefit from technological change. In particular, we seek to contribute to the strong predictions in parts of the literature of job losses and deskilling through automation and machine learning (Frey and Osborne, 2017).

Our evaluation criteria are motivated by the following three key observations that we find in the more recent literature: 1) A significant body of research forecasts negative labor market outcomes via risk indices, typically measuring the extent to which workers in an occupation can be replaced by technological solutions (Frey and Osborne, 2017). This prominently discussed job *destruction* hypothesis was empirically coinciding with *rising* employment levels. OECD labor force participation is at an all-time high and economies with considerable shares of *automatable* jobs, such as France, Ireland, or Slovakia, even show the steepest increases in employment levels (Georgieff and Milanez, 2021). By tracking research design and empirical operationalization, we compare this camp of evidence against other research, and assess to what extent evidence is a function of research design choice. 2) Automation is associated with rising employment levels in some countries, but net effects mask important channels of technological change that are evident when workers adapt their tasks on the job or retrain in-house (Acemoglu and Restrepo, 2022; Dauth et al., 2021), when workers need to change their employer, drop out of the labor force, or retire early (Cortes et al., 2017), and when workers select into previously unknown occupations (Autor et al., 2021). The literature has highlighted distributional effects and that automation comes with strong bias, either augmenting and supplementing or displacing and labor destructing certain labor market segments. Thus, we track existing evidence fourfold, according to whether automation affects low, middle, and highly skilled labor market segments, as well as the economy as a whole. 3) The speed of technological adaption as well as its geographical diffusion heavily depend on country-specific circumstances, and frontier technology is usually employed by highly productive firms in the most advanced economies. The same holds true for technology-induced labor market outcomes that are nested in political economies and mediated by institutions and welfare regulation. In our review, we emphasize the context factors of technological change. Understanding the role of institutional cross-country variation and regulation is key to ultimately derive recommendations and design policy options. We do this by assessing the geographical provenience of existing research and by reviewing empirical specifications according to whether they incorporate political economy factors.

Our exploratory scoping review is based on a sample of 197 key studies that are the outcome of PRISMA selection criteria. Section 2 explains inclusion and exclusion decisions, lays out coding strategies vis-a-vis computed variables, and references key contributions in order to justify selection decisions. Result section 3 highlights the findings of our scoping exercise, with conclusive evidence for high (middle) skilled segments with positive (negative) labor market outcomes, but fairly mixed evidence for low-skilled segments, and the economy as a whole. Next, we illustrate the evolution of two decades of technology-induced labor market research, and track research designs and fitted data sources. We identify a considerable share of studies (around 25 percent) with evidence exclusively based on occupational classifications and show that these contributions are significantly over-represented in the strand of research forecasting negative labor market outcomes. Lastly, we report on the geographical spread of the review sample. While comparative research sufficiently covers a good share of OECD countries, single-country research is predominantly concentrated in some economies, with the by far highest share of studies focusing on the US labour market. We finally discuss apparent research gaps with regard to institutional context factors, usually at the core interest of (comparative) political economy research. Section 4 discusses the limitations and implications of this scoping review as well as two main avenues for future research: The first one calling for a better and refined understanding of employment outcome differences between low and middle skilled labour market segments, especially beyond the US. Secondly, we call for more nuanced accounts of institutional factors that mediate *endogenous* technology adaption. The empirical interface between technological labor market change and institutional context bridges the gap to policy research to derive options for legislative action and discusses impacts of potential policy intervention.

2 Selection Procedure and Research Framework

Following Arksey and O'Malley (2005), this scoping review applies a five-stage framework, adopting a rigorous process of transparency, allowing replication of the search strategy, and boosting the reliability of the research findings: (1) establishing the initial research question and operationalization, (2) locating relevant studies, (3) study selection, (4) charting the data, and (5) compiling, summarizing, and reporting the results.

Our review includes contributions from different academic fields, and we spare an in-depth discussion of theory and potential causal mechanisms. Two economic standard models have, however, influenced most of the empirical literature, and guided the research framework of this scoping review, too. The education race literature posits that to keep up with the evolving demand of technology for more educated workers, countries continue to raise educational attainment. Skill-bias here means that new technologies require workers to undertake increasingly sophisticated and skill-intensive tasks that are associated with higher education degrees (Autor, 2022). Secondly, the task-polarization literature conceptualizes the process of work as a succession of tasks and determines which tasks will be performed by machines and which by humans. The possibility of *replacement* occurs here for those regularly codifiable actions that can be fully specified by a set of rules and processes, encoded in software, and performed by machines (a comprehensive overview is provided by Autor, 2022, who introduces the canon of economic literature dealing with technology-induced labor market implications).

2.1 Study Inclusion

We apply research standards following the Preferred Reporting Items for Systematic Reviews and Meta-analyses (PRISMA). On August 20, 2021, we searched ProQuest for published peer reviewed studies. On August 21, 2021, we searched Ebsco and Google Scholar, and included working papers and unpublished manuscripts as well. We restricted the search to articles available in English from January 1, 2002 with empirical focus on one or more OECD countries. The boolean search strategies are available in Annex 1. Studies were included in the review if they, (1) assessed the consequences of technological implementation in terms of labor outcomes (e.g. employment stability, wages), (2) were observational, quasi-experimental, or mathematical modelling studies, (3) exhibit at least one empirical, or reproducible finding, and (4) measured any employment or wage related outcome. Studies were excluded that (1) did not contain an empirical link between technological adoption and labor outcomes (for example, if they explored only general labor market dynamics), (2) cover a non-OECD country.

Figure 1 presents the process of selection and sorting criteria. For screening purposes, we employed a literature screening tool (DistillerSR) that removed 1986 duplicates from our initial stock of 4339 research papers. A first screening phase excluded studies based on their title and abstracts; the second screening reviewed 402 papers full-text, discarding those not in line with our above mentioned inclusion criteria. The final sample consists of 197 eligible studies, for each of which we document a battery of standard indicators: (1) bibliographic information, such as authorship, year of publication, journal, abstract, and URL; (2) descriptive information, such as publication type, data sources, countries under study, time frames of empirical specifications, policy recommendations, and research gaps identified.

112 peer-reviewed journal articles build the majority of included studies, followed by 33 working papers, 42 contributions from multilateral organisations, 5 book chapters, and 5 conference papers. 13 reports from the OECD (Organisation for Economic Co-operation and Development) make up the biggest share of organisational research. Regarding institutional working papers, reports from the IAB (Institut für Arbeitsmarkt und Berufsforschung) with 6, IZA (Institut zur Zukunft der Arbeit) with 5, NBER (The National Bureau of Economic Research) with 4 studies were the organisations providing the highest number of included studies.

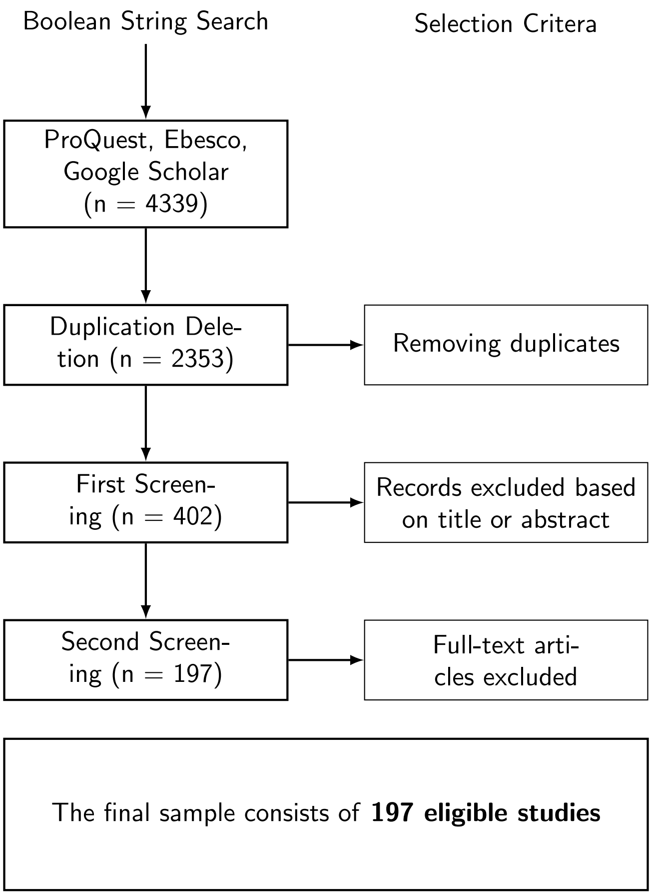

Figure 1: Prisma Chart and Study Selection

2.2 Heterogeneous Labor Market Outcomes

The main purpose of our review is to map and sort technology-induced labor market effects for different skill groups, following the task polarization literature (Acemoglu and Autor, 2011). We differentiate three different skill brackets (low, middle, high) to capture heterogeneous labor market outcomes and add a fourth outcome category for the net effect (i.e. the overall effect for a given economy). Skill-bias has been at the forefront of the related literature, but is operationalized in a variety of indicators, such as wage percentiles and occupational average pay, experience and tenure brackets, or educational degrees. Thus, we adopt the operationalization of a given study based on judgement calls. Next, we code outcome categories according to whether the study reports positive, neutral, or negative findings, which results in a 3 by 4 matrix for both employment effects and wage effects, as noted in Table 1. This coding scheme leads to an unbalanced data set, given that most studies focus on a subset of those 24 measures, as for instance a paper on employment effects leaves out a discussion of wage effects, or only exhibits results for some skill groups.

Table 1: A Review Matrix to Track Technological Change for Skill Groups

	Employment Effects			Wage Effects		
Treatment Group	Positive	Neutral	Negative	Positive	Neutral	Negative
Net Economy	0/1	0/1	0/1	0/1	0/1	0/1
High Skilled	0/1	0/1	0/1	0/1	0/1	0/1
Middle Skilled	0/1	0/1	0/1	0/1	0/1	0/1
Low Skilled	0/1	0/1	0/1	0/1	0/1	0/1

Notes: We apply this binary coding strategy to the baseline specification of each included study. Please note the following two caveats: 1) Identifying the core statistical model as the benchmark scenario is sometimes not easy, given the variety of specifications, or comparative research including more than one country. 2) Due to the different operationalizations of skill (e.g. as education, pay, tenure, occupations), coding relies on judgment calls of the reviewers, too.

2.3 Heterogeneous Technology Inputs

We distinguish three different automation technologies in line with existing research (Brynjolfsson et al., 2018; Felten et al., 2019; Webb, 2020): (1) Robots designed to automate manual repetitive tasks, (2) information and communication technology (ICT) and software solutions performing non-manual routine tasks, and (3) artificial intelligence applications that automate more advanced white-collar tasks, such as for for instance problem solving or logical reasoning.

2.4 Parametrization and Design Choice

Comparing findings of 197 papers is not trivial, considering different research designs, data sources, forms of technology, in different countries and time frames. In order to explain similar outcomes (i.e. employment and wages) with highly heterogeneous accounts of technological change, we develop a coding strategy that tracks four camps of operationalization strategies that we find prevalent in the literature: (1) The standard risk of automation approach, leveraging occupational classifications (Frey and Osborne, 2017; Goos et al., 2014), (2) the task-based approach, focusing on micro-compositions of work activities at the job level (Acemoglu and Autor, 2011; Autor, 2013; Arntz et al., 2017; Spitz-Oener, 2006), (3) an industry approach, highlighting sector-specific circumstances and investment variation (Klenert et al., 2020), and (4) a firm-based approach, stressing business effects, price elasticities, and competition dynamics (Acemoglu et al., 2020; Aghion et al., 2021; Bessen et al., 2020). By mapping those patterns in our sample, we can then assess to what extent those research designs explain different labor outcomes of technological change. We identify those four camps via data inputs/parametrization, and code them if econometric models contain (1) occupational classifications (Standard Occupational Classification (SOC), International Standard Classification of Occupations (ISCO), or else), (2) on the job task data (such as from O*NET, the Programme for the International Assessment of Adult Competencies (PIAAC), or BIBB/BAuA Employment Surveys), (2) industry or sector specific classifications (as for instance the North American Industry Classification System (NAICS), European classifications (NACE)), or (4) firm-based data or linked employee-employer data. Those four codes are not exclusive, as certain models leverage all four sorts of data.

3 Results

By setting the cutoff date to 2003, the publication year of an influential study by Autor et al. (2003), Figure 2 gives a first impression of our sample and plots the evolution of two decades of technology-induced labor market research. The sample incrementally increases, takes up significant gains in 2012 and reaches a preliminary peak in 2019 with 39 publications. As we include working papers and institutional research in our database query, the sample is biased towards more recent years, which can partly explain the relative thin coverage for the 2002-2010 period.

3.1 Parameters in the Research Field

What do data sources tell us about the evolution and advancements of the research field? The right panel of Figure 2 displays the above mentioned categorization of data inputs over time: Here we observe occupational classifications as highly prevalent parameters, which are applied in around 75 percent of all contributions, followed by industry classifications that are used in every second publication. Both levels of analysis are standard labor market research classifications, and usually build the backbone of econometric modelling. Occupational classifications take center stage in this research field, as they intuitively inform about the routine task intensity and indicate where investment in automation capital is most likely to be allocated. In the evolution of the literature, this potential for possible labor substitution (not *de facto* implementation) has attracted significant attention, and constitutes the well-known *risk of automation* literature (Acemoglu and Autor, 2011; Frey and Osborne, 2017; Goos et al., 2009). The same contributions published empirical indices that are standard practice in this field, and usually link routine/substitution scores to occupational classifications. These measures have been widely reproduced and belong to the standard tool-kit in the field.

An important refinement of the occupation argument comes from individual-level task applications, showing that tasks *within occupations* are susceptible to automation (Autor et al., 2003; Arntz et al., 2017; Böhm, 2020). They added considerable nuance to the job destruction hypothesis by for instance Frey and Osborne (2017), and substantially revised the overall number of jobs at risk of automation downwards. The spread of such research models is depicted as the green line in Figure 2, and represents up to 25 percent of research in our sample.

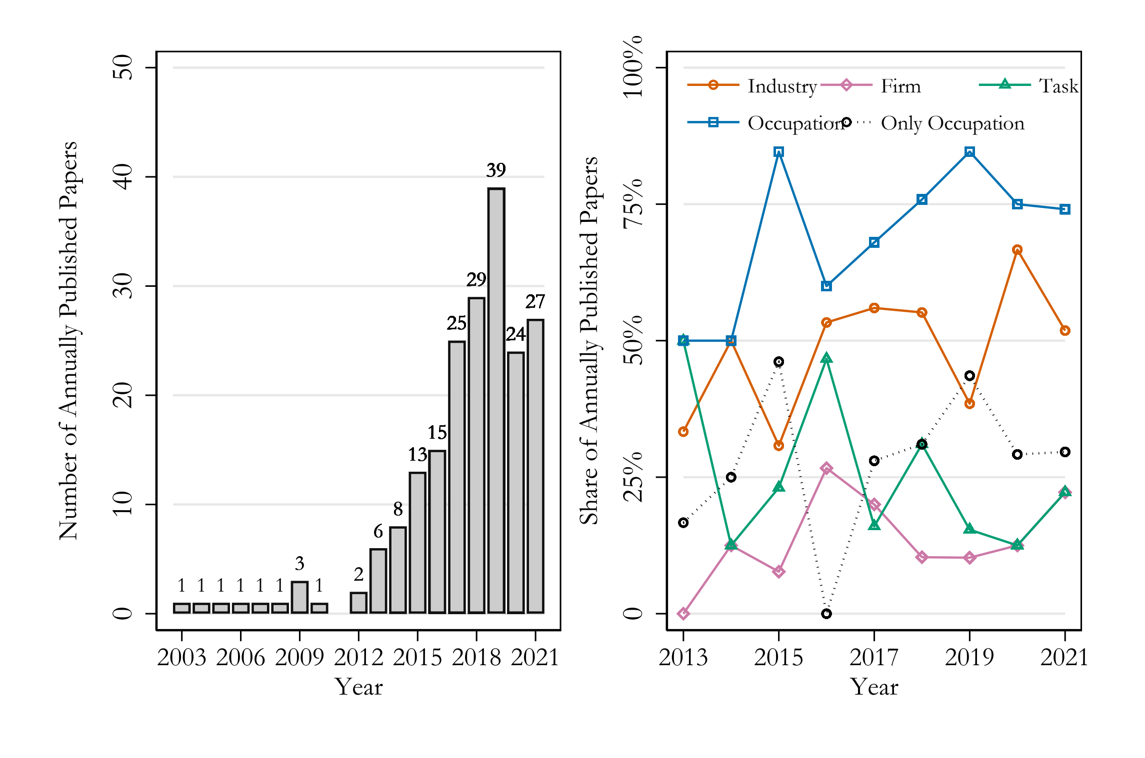

Figure 2: Identification of Parameters in the Literature

Notes: This graph shows absolute and relative frequencies of different parameters in 197 papers. Identified parameters in the right panel are not exclusive categories, with the exemption of the dotted line, tracking the percentage of studies that employ a single identification strategy via occupational classifications.

Alongside these two research camps –occupation based and task based–, we identify a third camp of research taking into account business dynamics and inter-firm competition, of what we coin firm-based indicators for the purpose of this review (for a recent discussion see Aghion et al., 2021). In this conjecture, business-stealing mechanisms make some firms more productive at the expense of others, hence automating firms are able to *increase* their headcount due to more demand and quality-adjusted prices, while competitors lose out and revise their staff downward. (Acemoglu et al., 2020; Aghion et al., 2021; Bessen et al., 2020). In our sample, we observe that such data sources are used in every 10th study, with a slight increase in more recent years (the pink line).

The identified research patterns in Figure 2 are not exclusive. Task-based micro data can be well combined with employer data to observe how technology alters the skill requirements in firms (as for instance in Acemoglu et al., 2022). However, we find a substantial body of research solely applying occupation based measures, as indicated in the black dotted line in Figure 2.

3.2 Disaggregate Effects of Automation

How conclusive is the evidence of labor market research addressing the consequences of automation technology for workers' employment stability and pay? In Table 1, we present a heterogeneous review matrix, aimed to capture employment dynamics beyond the aggregate labor displacement hypothesis. Figure 3 displays the distribution of research findings differing by skill group. Looking at the impact on the economy as a whole, we find for both wages and employment levels fairly scattered evidence. For employment related outcomes, we count 19 papers with evidence of negative net-employment, while at the same time 29 with neutral results, and 25 with positive evidence; a similar distribution, although with overall smaller numbers, occurs for wage effects on economy level, where 6 papers observe negative net-employment, 7 papers report neutral evidence, and 10 papers document positive outcomes for overall employment levels.

For the three following skill categories, we observe roughly the view espoused by the task polarization literature, with the largest share of positive evidence attributed to high-skilled workers. Among 86 papers with evidence for middle-skilled workers (86 percent), 74 report negative outcomes for employment levels, and the share in the wage category is only marginally lower (75 percent). For low-skilled workers, our results again present a mixed picture: 37 studies find negative employment outcomes, 10 neutral, and 39 positive effects.

The fact that the evidence regarding employment and wages for high-skilled and middle-skilled workers is relatively uncontested (i.e. around 80 percent or more in the respective category) and moving strongly in countervailing directions provides conclusive evidence for the task polarization model à la Autor (2022). According to the same literature, low-skilled labor segments are supposed to experience employment gains as well (the other part of the U-shaped employment polarization). Even though our framework is not per-se designed to rule out specific mechanisms, we can report, however, that the evidence in this skill bracket is far from being conclusive.

Mapping existing evidence within the framework presented, our sample presents a rather concise picture for high and middle skilled labor segments, while evidence for overall employment on the macro level of the economy as well as for low-skilled workers remains mixed. We recognize the danger of comparing studies that vary considerably in terms of automation, country, time and research design. We are also aware of publication bias, likely affecting our sample by studies with either negative or positive outcomes with statistical significance. However, we contend that it is unlikely that these biases affect the four categories in Figure 3 differently. Please also note that Appendix A contains the same figure of labour market effects differentiated by the sort of technology.

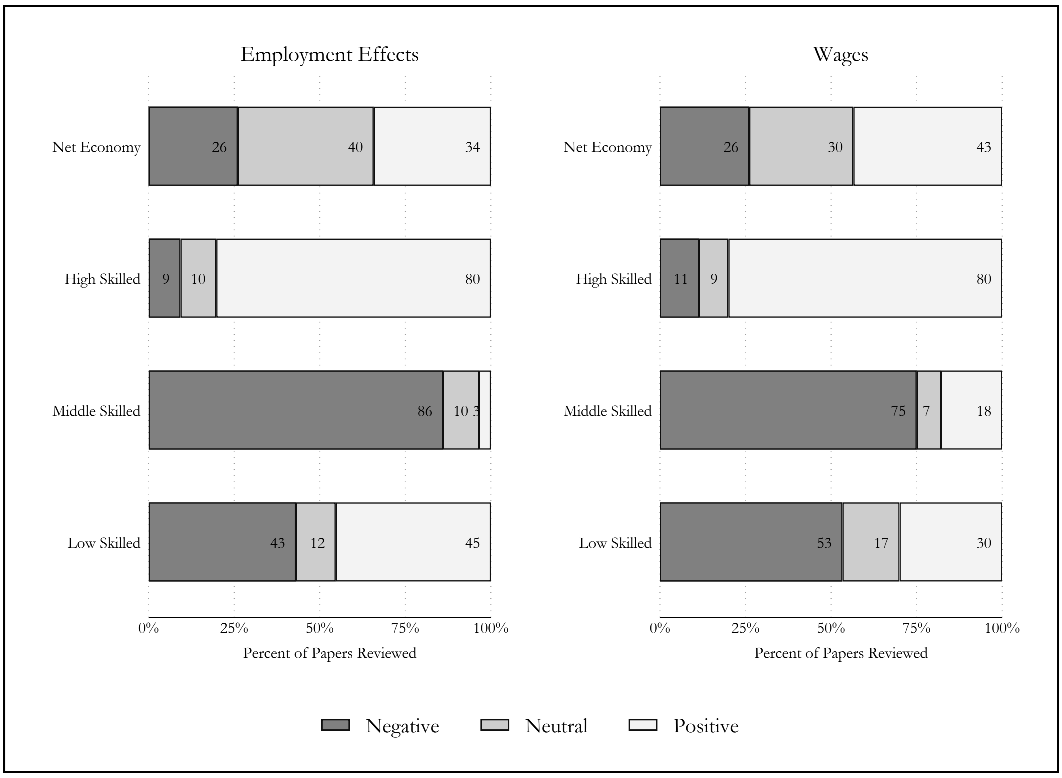

Figure 3: Heterogeneous Effects of Labor Market Automation

Notes: Relative shares of papers indicating labor market outcomes for skill groups. Please note that absolute numbers differ by category: For the net-economy category, we reviewed 19 with negative, 29 with neutral, and 25 with positive employment effects, hereafter annotated as (19/29/25). In the high-skilled bracket: (8/9/69); middle-skilled: (74/9/3); low-skilled: (37/10/39). With regard to wages, we count for the net-economy: (6/7/10); the high-skilled: (4/3/28); the middle-skilled: (21/2/5); low-skilled: (16/5/9).

3.3 Research Design Choice

What factors contribute to explaining the variation in the reported results? Above, we gave a qualitative, non-exhaustive overview of evidence by stressing important data and study dimensions, and Figure 2 illustrated these distinct camps by tracking study parameters over time. We now try to establish an empirical link between the ambiguous evidence for labor market outcomes and research parametrization. We run a series of simple regression models to test whether data choice correlates with negative or positive evidence for the impact of technological change on the labor market. For each included study, we constructed a dependent variable that sums up positive/negative outcomes and omits neutral values. For instance, a paper that finds negative employment effects on all four heterogeneous levels is coded as

4, while automatically shows the value 0 for the regression specification for positive employment effects. This coding strategy creates the maximum variation, to assess if research designs tend to correlative with rather negative or positive labor market outcomes. As independent variables, we introduce eight different dummy variables, consisting of the five different data parameters (firm, occupational, industry, task, and occupational data only) plus binary information about the sort of technology. Before turning to the multivariate regression results, we highlight several caveats that considerably limit the scope of our findings. Instead of coding estimated standardised coefficients from included studies, we code outcome categories as positive, neutral, or negative, thus our method does not meet common standard requirements for meta-regression analyses, and other methods usually applied in systematic reviews. Second, even though we take into account non-findings by coding results as neutral, we are not able to empirically assess the extent of publication bias, which is likely affecting our sample of included studies.

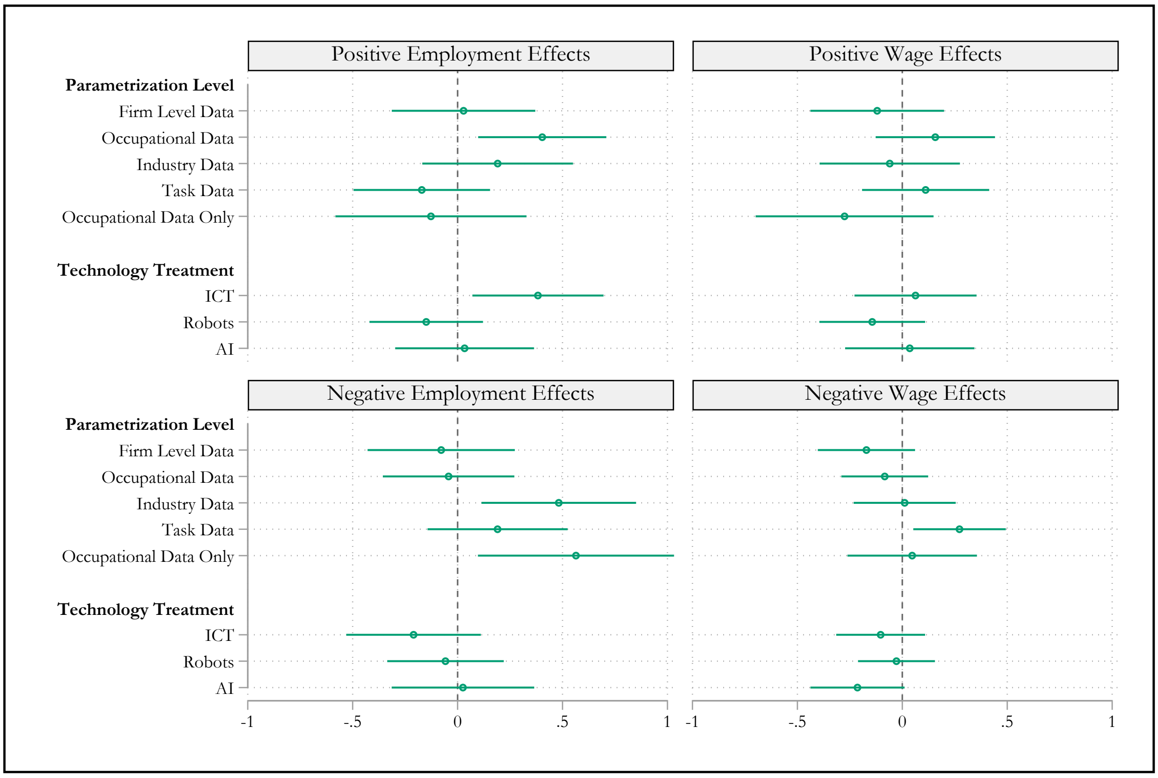

Figure 4: Average Marginal Effects of Parametrization on Outcome Variables

Notes: All four panels show average marginal effects of four linear regression models. In Panel A, the dependent variable is computed as the sum of positive employment effects, ranging from 0 to a maximum of 4 (positive effects for net-economy, high-skilled, middle-skilled, low-skilled brackets). Panel B, computes the sum of positive wage effects in the same vein. Panel C and D compute the sum of negative effects, ranging from 0 to a maximum of 4 with all four segments showing negative outcomes).

Figure 4 shows average marginal effects for four different regression specifications, and illustrates the eight main factors that have been discussed in the previous section and that we think are important in terms of having an impact on the effect of technology on labor outcomes. Some central robust correlations stand out: Panel A shows that studies including ICT data tend to find on average higher positive employment effects across all four categories. Our technology dummy variables are not mutually exhaustive, therefore we cannot directly asses if information and communication technology itself correlates with positive employment effects, (important papers include ICT data in addition to other sources, as for instance Graetz and Michaels, 2018). Turning to Panel C, which focuses on the evidence of negative employment effects, we find a robust correlation between negative employment reporting and the use of industry parametrization and occupational data. Here, we expect the camp of risk of automation literature to drive variation, and we find that the evidence of studies exclusively leveraging occupational classifications report on average more negative employment outcomes. Given that this significant body of literature focuses on how susceptible jobs are (i.e. this sub-field is biased towards negative outcomes) this result is not too surprising. However, our results show that this is not the case if studies include occupational data in combination with other parameters. Thus, we can conclude that using occupational data in combination with other sources is a clear quality criterion and should be an imperative for future research. Lastly, Panel D shows that individual-level task data correlates with more negative wage effects reported. Notably, the number of papers in this bracket is rather small, and likely driven by the contributions of Bessen et al. (2019); Dauth et al. (2021).

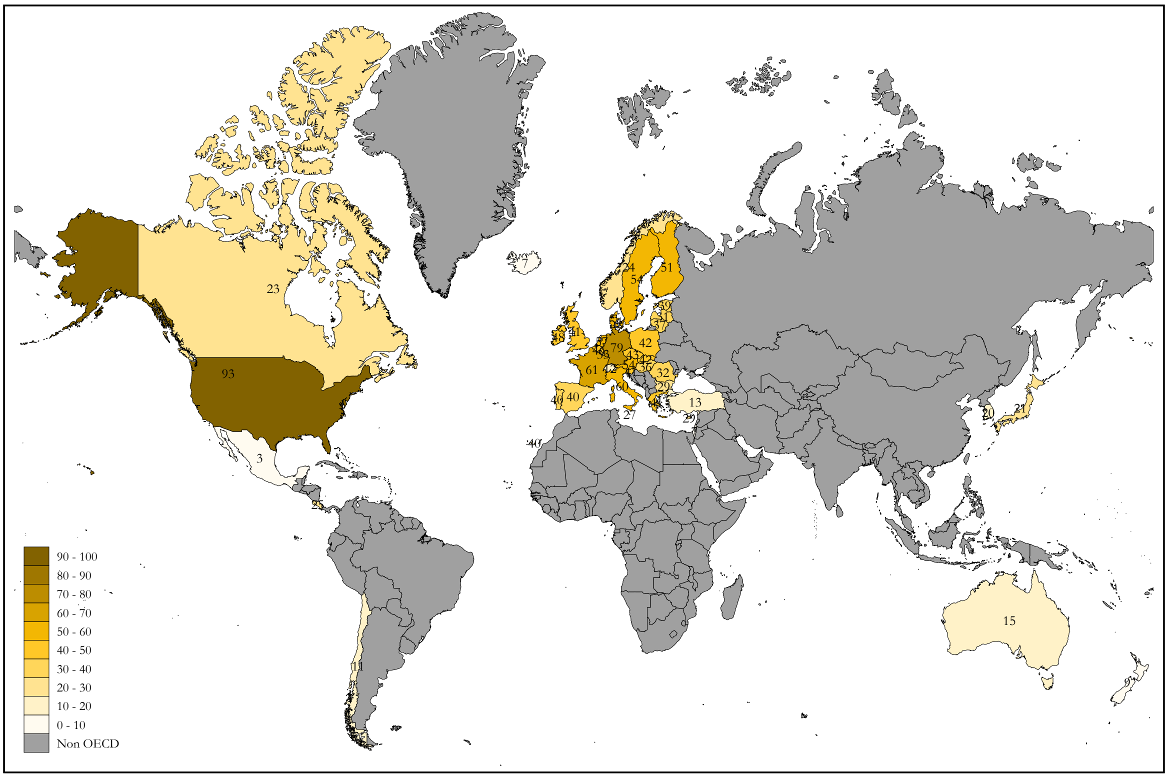

Figure 5: Geographical Coverage of Included Studies

Notes: This map counts the overall spread of research across OECD countries. A significant share of included studies is comparative research, leveraging international data, such as the Programme for the International Assessment of Adult Competencies (PIAAC), and therefore covering a large number of countries.

3.4 Geographical Spread and Comparative Political Economy

In this section, we map the geographical coverage of our sample, and show that only a small fraction of existing research provides context factors of political economy research. Starting with coverage, Figure 5 depicts the overall spread of our review sample across the OECD. Out of 197 studies, almost every second study (93 in total) covers the US, followed by 79 studies for Germany, and 61 for France. The least covered countries are Mexico (3), Costa Rica (2), and Chile (11). Even though we find considerable variation in terms of geographical focus, European countries are covered with an average of around 40 papers. This is partly due to research designs using international comparative data sets for groups of several countries. The OECD's Programme for the International Assessment of Adult Competencies (PIAAC), for instance, covers most OECD countries, and represents a widely used data set in the field (e.g. Nedelkoska and Quintini, 2018).

Using data variation across countries, however, can be problematic if institutional factors are vaguely specified, implicitly supporting the assumption that automation uniformly leads to ceteris paribus labor outcomes, regardless of varying macroeconomic context, or national institutions, such as welfare and social safety nets, or skills and training schemes. For instance, a recent OECD report shows that employment growth is most evident in exactly those countries, whose labor force has been estimated to be subject to considerable risk of automation, such as France, Ireland, or Slovakia (Georgieff and Milanez, 2021). Identification strategies that spare institutional context (i.e. especially cross-country designs) are therefore ill-suited to track the adaption of labor to automation.

What do we know about institutional factors mediating the effect of technological change on the labor market? To provide an overview over existing research we limited our analysis to studies with a focus on a single political economy, given that there is no comparative study taking into account labor market institutions. Figure 6 counts 127 studies using intra-national designs, i.e. with a focus on automation implications for a single political economy. Again, single country studies are highly concentrated with a strong focus on the US labor marked (62 studies), following a significant gap to the German labor market (20 studies). Among these studies, we report a considerable lack of research agendas addressing labor market institutions. There is abundant evidence documenting the role of labor market institutions in shaping employment stability and wage outcomes (as for instance via labor representation), but out of 197 included studies we are only aware of two contributions addressing such institutional context (Dauth et al., 2021; Parolin, 2021). In the US context, Parolin (2021) shows that in regions and industries with high unionization rates, wages of highly routinized occupations remained stable, but overall employment shares in these occupations significantly dropped. Supporting evidence comes from Dauth et al. (2021), who find find that robot-induced labor displacement was cushioned by regional unionization in Germany.

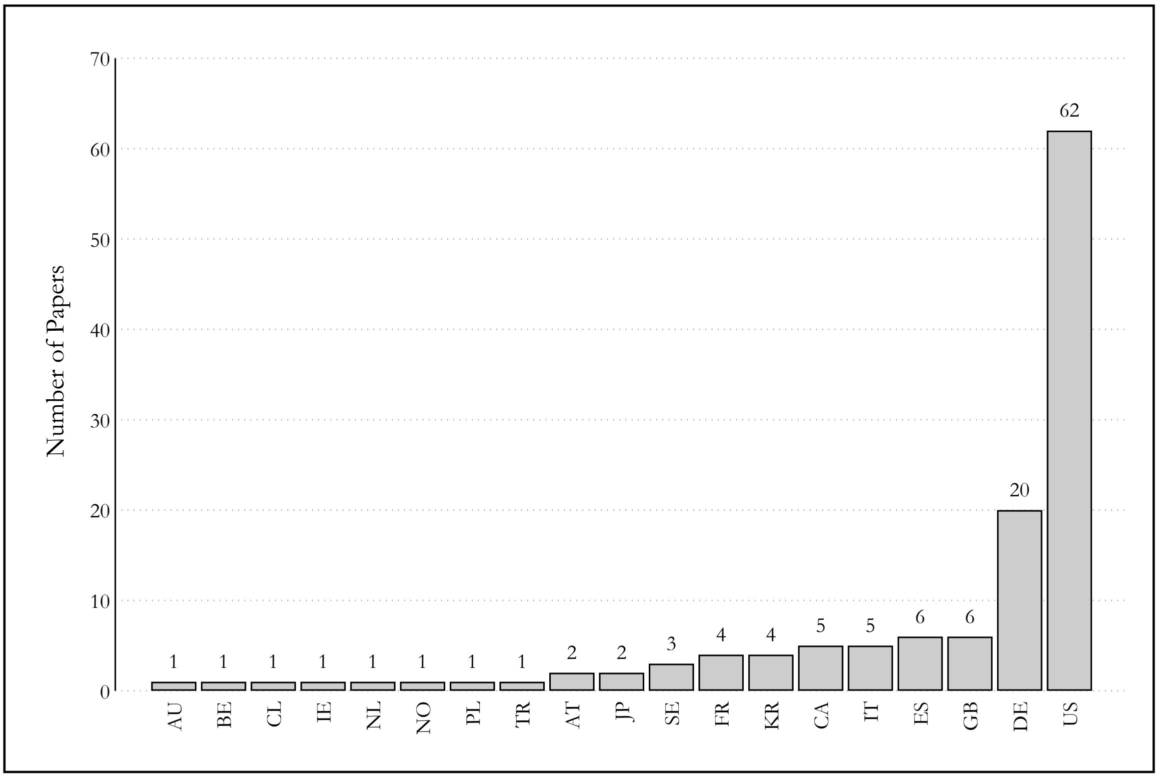

Figure 6: Single Country Studies are Concentrated in A Small Fraction of Countries

Notes: Single country research as included studies with focus on one country or specific regions of a single country.

There are other important channels beyond labor representation via unions that will likely affect technology's impact on employment and wages. In Germany, shop-floor representation in the form of works councils are part of decision-making processes in management, and have been found to shape technological implementation strategies at the firm level (Genz et al., 2019). Collective Bargaining Agreements are typically negotiated at region/industry level and set wage levels, working conditions, as well as training schemes for a large part of the German workforce. Overall, the evidence base that links institutions to automation driven labor market outcomes is thin or even non-existent for the majority of countries beyond the US and Germany. Drawing from industrial relations research is not only important to better understand the ways labor markets adapt to technological advancements. Institutions are constituting the important interface for which policy makers can then fine-tune adjustment programmes. Exploring such avenues of endogenous technological change remains crucial to.

4 Conclusion

What is known from the available literature about the impact of automation technologies on wages and employment levels across the OECD? Using a scoping review framework, we structured the existent literature according to (1) automation type, (2) empirical modelling strategies, and (3) heterogeneous effects for skill groups. The findings can be summarized in three key take-away's: We find conclusive evidence for high (middle) skilled segments with positive (negative) labor market outcomes, but fairly mixed evidence for low-skilled segments, and the economy as a whole. Contrasting the strong emphasis on labour market *polarization*, employment outcomes for low-skilled work remain an empirically open question, especially beyond the US context. Therefore, future research should take a clear focus on the task content of low-skilled work, and to what extent it differs from middle-skilled, routine work. Second, by illustrating underlying research designs and fitted data sources, we identify a considerable share of studies (around 25 percent) with evidence exclusively based on occupational classifications and show that these contributions are significantly overrepresented in the strand of research forecasting negative labor market outcomes. Thirdly, we report on the geographical spread of the review sample. While comparative research sufficiently covers a good share of OECD countries, single-country research is predominantly concentrated in some economies, with the by far highest share of studies focusing on the US economy. Existing research has not sufficiently explored institutional factors that shape both technology implementation, as well as the implications for the work force. Future studies should continue to identify labour market institutions that mediate the technology-work nexus, and take up research to compare country settings. Understanding the relationship between technological labor market change and institutional context is crucial for policy research to derive options for legislative action and discusses impacts of potential policy intervention.

This scoping review comes with one limitation: Compared to systematic reviews and meta-regression methods, this review does not systematically report on all empirical specifications and estimates of included studies. Given the bulk of comparative studies reporting different country estimates within the same specification, future reviews could assess effects sizes more thoroughly and analyse the heterogeneity of reported estimates.

References

Acemoglu, D. and Autor, D. (2011). *Skills, tasks and technologies: Implications for employment and earnings*, volume 4.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Ai 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. (2022). Tasks, automation, and the rise in us wage inequality. *Econometrica*.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. (2021). The direct and indirect effects of automation on employment: A survey of the recent literature. *Työpaperi. Saatavilla: https://scholar. harvard. edu/files/aghion/files/direct_ and_ indirect_ effec ts_ of_ automation. pdf, viitattu*, 23:2022.

Arksey, H. and O'Malley, L. (2005). Scoping studies: towards a methodological framework. *International journal of social research methodology*, 8(1):19–32.

Arntz, M., Gregory, T., and Zierahn, U. (2017). Revisiting the risk of automation. *Economics Letters*, 159:157–160.

Autor, D. (2013). The "task approach" to labor markets: an overview. *Journal for Labour Market Research*, 46(3):185–199.

Autor, D. (2022). The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty. *National Bureau of Economic Research NBER Working Paper 30074*.

Autor, D., Salomons, A., and Seegmiller, B. (2021). New frontiers: The origins and content of new work, 1940–2018. *MIT Working Paper, July*.

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

Bessen, J., Goos, M., Salomons, A., and Van den Berge, W. (2019). Automatic reaction-what happens to workers at firms that automate?

Bessen, J., Goos, M., Salomons, A., and van den Berge, W. (2020). Firm-level automation: Evidence from the netherlands. *AEA Papers and Proceedings*, 110:389–93.

Böhm, M. J. (2020). The price of polarization: Estimating task prices under routine-biased technical change. *Quantitative Economics*, 11(2):761–799.

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.

Cortes, G. M., Jaimovich, N., and Siu, H. E. (2017). Disappearing routine jobs: Who, how, and why? pages 69–87.

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

Felten, E., Raj, M., and Seamans, R. (2019). The occupational impact of artificial intelligence: Labor, skills, and polarization: The role of complementary skills and technologies",.

Frey, C. B. and Osborne, M. A. (2017). The future of employment: How susceptible are jobs to computerisation? *Technological Forecasting and Social Change*, 114:254–280.

Genz, S., Bellmann, L., and Matthes, B. (2019). Do german works councils counter or foster the implementation of digital technologies? *Jahrbücher für Nationalökonomie und Statistik*, 239(3):523–564.

Georgieff, A. and Milanez, A. (2021). What happened to jobs at high risk of automation? *OECD Social, Employment and Migration Working Papers, No. 255.*

Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in europe. *Source: The American Economic Review*, 99:58–63.

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

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

Klenert, D., Fernández-Macías, E., and Antón, J.-I. (2020). Do robots really destroy jobs? evidence from europe jrc working papers series on labour, education and technology 2020/01.

Nedelkoska, L. and Quintini, G. (2018). Automation, skills use and training. *OECD ELS Working Paper*.

Parolin, Z. (2021). Automation, occupational earnings trends, and the moderating role of organized labor. *Social Forces*, 99(3):921–946.

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

Webb, M. (2020). The impact of artificial intelligence on the labor market. *SSRN: https://papers.ssrn.com/sol3/papers.cfm?abstract*$_i$*d* = 3482150.

A Appendix

1. **Search string 1**: TX all text (automation digitalisation robots computer) AND TX all text (jobs workers workforce employment labour unemployment polarisation inequality) AND (wages salary income pay earnings skill);

2. **Search String 2**: TX all text (Job Polarization, Task-Biased, routine, automation);

3. **Search String 3**: job polarisation, routine, automation, employment labour robots.

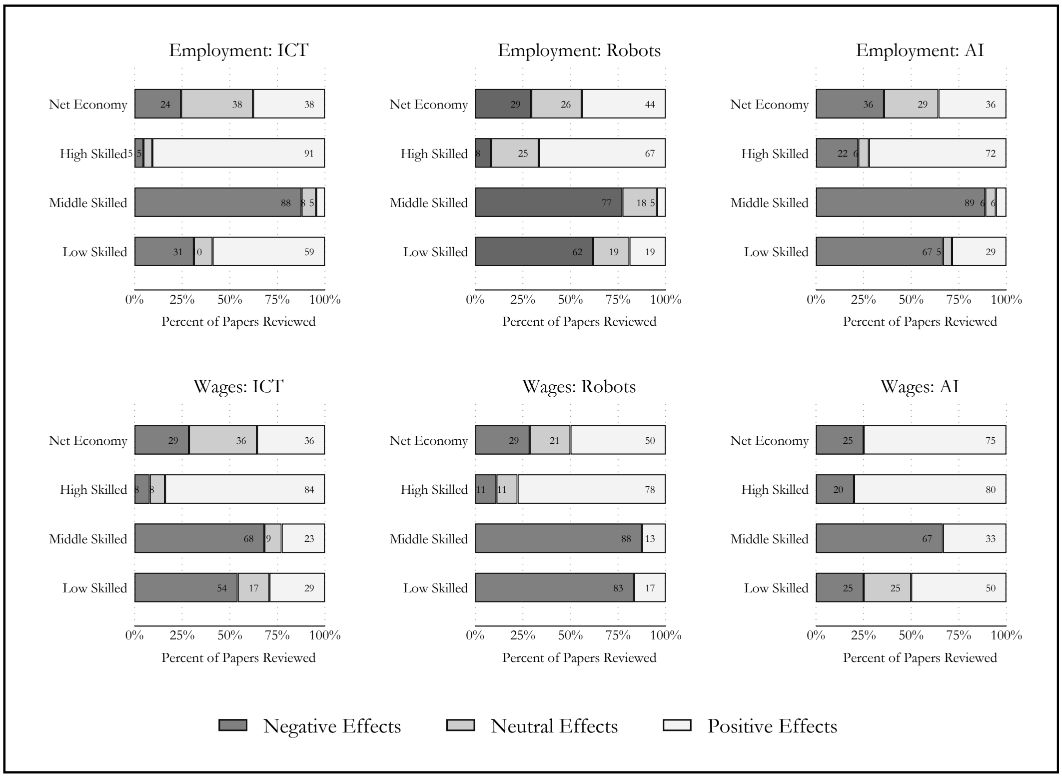

Figure 7: Employment and Wage Effects By Type of Automation

Notes: Relative shares of papers indicating labor market outcomes for skill groups. Please note that absolute numbers differ by category. The employment panels come with the following absolute papers: ICT: Net Economy (11/17/17), High Skilled (3/3/58), Middle Skilled (57/5/3), Low Skilled (19/6/36). Robots: Net Economy (10/9/15), High Skilled (2/6/16), Middle Skilled (17/4/1), Low Skilled (13/4/4). AI: Net Economy (5/4/5), High Skilled (4/1/13), Middle Skilled (16/1/1), Low Skilled (14/1/6). The wage panels base on the following absolute numbers: ICT: Net Economy (11/17/17), High Skilled (3/3/58), Middle Skilled (57/5/3), Low Skilled (19/6/36). Robots: Net Economy (10/9/15), High Skilled (2/6/16), Middle Skilled (17/4/1), Low Skilled (13/4/4). AI: Net Economy (5/4/5), High Skilled (4/1/13), Middle Skilled (16/1/1), Low Skilled (14/1/6).