Technological Change and Labour Market Institutions and Their Effect on Employment, Wages, and Inequality

Matthias Haslberger

Nuffield College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy Social Policy

Trinity 2021

To my father Konrad

Acknowledgements

I have incurred numerous debts in writing this doctoral thesis. Through the ups and downs of this 4-year journey, I have benefited from the guidance, advice, and support of many individuals more knowledgeable and accomplished than myself.

First of all, I want to thank my supervisors, Erzsebet Bukodi and Brian Nolan, for striking the right balance between encouraging me to explore my academic interests and demanding deliverable results, and for supporting me in numerous ways with feedback, reference letters, and advice. Bess, who was also my Masters supervisor, is a main reason why I decided to pursue a DPhil in the first place. Her methodological advice and encouragement to push on when I felt stuck or low on motivation, was invaluable. So was the feedback of Brian, whose emphasis on clarity forced me to sharpen my thinking and writing so that the end product, I hope, makes for an enjoyable read.

Richard Breen, too, has been a great help as my college supervisor and I am grateful for his constructive feedback and pastoral support. I also want to thank Jane Gingrich for the work with SCHOOLPOL which has been a stimulating change from my thesis research and a crucial source of funding for my 4$^{\text{th}}$DPhil year. I have also benefited a great deal from the comments of my assessors for the Transfer and Confirmation of Status, Bernhard Ebbinghaus and Craig Holmes.

In my efforts to navigate the world of academic publishing and the job market, I have at various points benefited from the advice and example, as well as critical feedback, of fellow early career researchers, including Marta Golin, Tobias Rüttenauer, David Weisstanner, and Dirk Witteveen.

I am also grateful to the Department of Social Policy and Intervention and Nuffield College for their generous financial support. Without the Barnett House-Nuffield scholarship, I could not have pursued my passion for academic research in the first place, and their numerous travel grants allowed me to share my research with others and build important connections.

My thanks also go to the staff at Nuffield College. The periods of college closure during the pandemic showed what an immense difference it makes for one's academic work to have all the little worries taken care of. Buttery, cleaning, IT, kitchen, library, lodge, and maintenance staff all have been truly exceptional at making my life as a student easier and rarely receive the credit they deserve.

I am also indebted to the wonderful communities which I had the privilege to be a part of. In Nuffield, I am thankful above all to my dear friends Mattia Bertazzini, Sarah Clifford, Artur Doshchyn, Marta Golin, Said Hassan, Tuuli-Anna Huikuri, Fijnanda van Klingeren, Max Marczinek, Tanisha Mohapatra, Tobias Rüttenauer, Ahmed Tohamy, Florianne Verkroost, and Yuan Yi Zhu, as well as the Old Guard including Jasmine Bhatia, Noah Carl, Roberto Cerina, and Arshia Hashemi, for the bike rides, dinners, concerts, walks, and discussions about life, politics, philosophy (and sometimes utter nonsense). Outside College, Marilena Anastasopoulou, Sharika Khan, Juhi Kore, and Harshmeena Sanghani have been my biggest cheerleaders and an unfailing source of support. Raphael Heim and Ertugrul Polat, who turned from housemates into good friends, and my friends in Linacre Boat Club and Christ Church Boat Club, above all Brad Barlow, Mehran Hosseini, Derek Xu, and Pete Young, have helped me maintain a healthy balance between work and private life. To them, and numerous others, I shall forever be grateful for making Oxford a place that is home to many of my most cherished memories.

At last, my deepest thanks to my family. To my mother Christine whose unwavering support and profound trust in my ability to see through a project of this proportion enabled me to take this journey. To my brother Martin, who, always hot on my heels, has been a source of friendly competition and many a wise word. Finally, to my father Konrad, who did not live to see me embark on this journey, yet whose example of curiosity and dedication has ultimately been my greatest source of inspiration. This thesis is dedicated to him.

Abstract

Technological change and labour market institutions (LMI) exert a major influence on the labour markets of developed economies. Routine-biased technological change (RBTC) has been linked to employment polarisation in numerous countries, while institutions such as unions and employment protections have been found to reduce inequalities between workers. However, much is still unknown about the labour market impact of technological change across countries and the interplay between technology and institutions. In this thesis, I integrate insights from labour economics and sociology to address some of the gaps in this literature. I investigate three closely linked research questions: how is routine intensity best measured and how does it vary between countries and over time? Under which conditions does RBTC lead to employment polarisation? And do robotisation and LMI contribute to the patterns that enable polarisation? Using data from the European Working Conditions Survey, the Luxembourg Income Study, and the International Federation of Robotics, I pursue these questions in a sample of OECD countries during the period from 1993 until 2016. First, I develop improved indices of occupational task content which show meaningful differences with established measures as well as across countries and over time. Secondly, I demonstrate that RBTC is associated with employment polarisation only in countries where routine occupations are concentrated in the middle of the wage hierarchy. Finally, my analyses suggest that robotisation has reduced the relative wages of routine manufacturing occupations and thus made employment polarisation less likely, particularly in countries with strong employment protections for temporary workers. The findings of this thesis advance the study of employment and wage changes in both economics and sociology, and illustrate the benefits of jointly analysing technological and institutional explanations for differences and changes in employment, wages, and inequality.

Contents

List of Figures

List of Tables

List of Abbreviations

AI . . . . artificial intelligence

CPDS . . . . . Comparative Political Data Set

CPS . . . . . . Current Population Survey

DOT . . . . . . Dictionary of Occupational Titles

EPL . . . . . . employment protection legislation

EU-LFS . . . . EU Labour Force Survey

EU-SES . . . . EU Structure of Earnings Survey

EU-SILC . . . EU Statistics on Income and Living Conditions

EWCS . . . . . European Working Conditions Survey

FYFT . . . . . full-year, full-time

GSOEP . . . . German Socioeconomic Panel

IABS . . . . . Institute for Employment Research Employment Samples

ICT . . . . . . information and communication technology

IFR . . . . . . International Federation of Robotics

ISCO . . . . . International Standard Classification of Occupations

ISEI . . . . . . International Socio-Economic Index of Occupational Status

ISIC . . . . . . International Standard Industrial Classification of All Economic Activities

IT . . . . information technology

LFP . . . . . . labour force participation

LIS . . . . Luxembourg Income Study

LMI . . . . . . labour market institutions

MW . . . . . . minimum wage

NACE . . . . . Statistical Classification of Economic Activities in the European Community

OECD . . . . . Organisation for Economic Co-operation and Development

O*NET . . . . Occupational Information Network

PPP . . . . . . purchasing power parity

RBTC . . . . . routine-biased technological change

RTI . . . . . . routine task intensity

SBTC . . . . . skill-biased technological change

SIOPS . . . . . Standard International Occupational Prestige Scale

USD . . . . . . United States Dollar

1 Introduction

Contents

1.1 The research puzzle and research questions

Labour market trends such as a polarising employment structure and widespread increases wage inequality have attracted considerable scholarly interest in recent years. This has led to the emergence of several strands of literature which attribute these trends to numerous factors, including technological change, the liberalisation of labour market institutions, demographic changes, and globalisation. Of these, technological change and labour market institutions are by many considered the most impactful. While economists have been the leading champions of theories of biased technological change, sociologists have more commonly advanced institutional approaches. Dialogue between proponents of the two approaches has been limited.

Despite this, theories of skill- and routine-biased technological change (RBTC) and worker bargaining power both provide compelling frameworks for the analysis of labour market trends since the 1990s. The fundamental importance of the explosion of computing power and the adoption of information and communication technology (ICT) for the structural changes that have occurred in advanced countries' labour markets is widely acknowledged. RBTC, as I explain in chapter 2, makes a set of specific predictions how this process should affect labour markets. Moreover, technological change, as arguably the most fundamental force driving the reorganisation of labour markets, should play out very similarly across developed countries and contribute to similar employment and wage trends. Finding, as I do below, that it does not, is indicative that technology may not be the full story, contrary to the impression that is fostered in some, though by no means all, of the economics literature.

Labour market institutions (LMI), because of their greater country differences, are the most obvious candidate for rationalising any differences in employment trends that biased technological change cannot explain. LMI have been comparatively neglected in the technological change literature, despite the vast parallel literature, including in economics, that documents the impact of unions and other institutions on employment and wages, as well as the erosion of inclusive LMI in many advanced economies. Thus, it seems that there is a great deal that can be learned by integrating these literatures and investigating the ways in which institutions may alter the way technological change plays out in the labour market. Greater attention to the interaction of institutional factors and technological change has often been demanded in the literature (Antonczyk, DeLeire & Fitzenberger 2018, Lemieux 2008), but it has so far not been implemented sufficiently.[1]

The overall research puzzle that I seek to illuminate in this thesis is therefore how technological change and labour market institutions affect employment, wages, and inequality. Investigating the impact of technological and institutional factors jointly, and integrating usually separate strands of literature from labour economics,

[1]The focus on technological change and LMI does not mean that I will disregard demographic changes or globalisation. Any meaningful analysis of technology and institutions must at least discuss these factors. However, the focus in this thesis is firmly on the interplay of technology and institutions.

economic sociology, social policy, and comparative politics represents a major contribution of this thesis.

The concrete research questions that I set out to answer in this thesis are the following:

1. How can occupational task content be measured in a way that corresponds to the underlying theoretical concepts and that accounts for the relevant dimensions of variation?

2. Can the diverse patterns of employment change in developed countries be reconciled with the key tenet of RBTC, that technological change everywhere substitutes for routine workers?

3. Do robotisation and LMI, independently and jointly, have heterogeneous effects on the wages of different occupational groups?

These research questions build on each other in a coherent logical sequence and allow me to investigate several interesting issues that are relevant to the labour market literature. My basic framework of analysis remains the RBTC model, which in my view still provides a compelling explanation for many of the empirical patterns that can be observed in developed countries. Indeed, as I shall argue, with the appropriate theoretical refinements, RBTC can account for some of the patterns that have previously invited criticism from sociologists and some economists.

The first research question thus addresses one of the fundamental weaknesses of the existing economic literature: its use of occupational task measures that are not always fit for purpose. Equipped with better measures, the second research question tackles the discrepancy between the pervasive polarisation following a decline of middle-wage routine jobs predicted by RBTC theory and the diverse patterns observed in reality. Finally, the third research question follows up on the key findings from the analysis of research question 2 and asks whether differences in robotisation and LMI, and interactions between the two, can explain the differences found there. Taken together, the answers to these research questions bring us closer to understanding the effect of technological change and labour market institutions on employment, wages, and inequality.

The research in this thesis advances knowledge in labour economics and economic sociology, as well as adjacent fields. Both economics and sociology are deeply invested in determining the effect of technological change on labour markets. However, the predominant model in labour economics does not account for institutions, and its unnecessarily stringent assumptions regarding the position of routine workers in the wage hierarchy limit its predictive usefulness.[2] Sociologists, while pointing out many of the flaws of the economic model, have sometimes gone too far and rejected RBTC altogether. With this thesis, I contribute to integrating the two disciplines that have done the most for our understanding of recent labour market changes.

My research is also relevant for the welfare state and comparative institutions literatures. For example, it shows the continued relevance of the power resources framework when analysing occupational wage hierarchies. The findings of this thesis suggest a renewed need for an analysis of the institutional determinants of inequality that seems to have lost some of its appeal since the early 2000s (Koeniger, Leonardi & Nunziata 2007, Pontusson, Rueda & Way 2002, Rueda & Pontusson 2000). Furthermore, chapter 7 of this thesis makes the case for paying greater attention to the small differences between countries' occupational hierarchies, their overall similarity notwithstanding (Hout & DiPrete 2006).

My contributions are both theoretical and empirical. On the one hand, I suggest refinements to the way in which occupational tasks are theorised and to the canonical version of the RBTC model. However, each of these theoretical contributions is accompanied by empirical evidence to underscore its relevance and robustness. These contributions are explicated in greater detail in chapter 3, since after thoroughly analysing the relevant literature, the value of the contributions of my thesis will be clearer. The remainder of this introductory chapter provides some

[2] This is not to say that no economists have analysed the impact of labour market institutions. For example, there is a vast amount of research in economics on the role of minimum wages and unions, some of which is discussed below. However, the RBTC literature in labour economics has developed with very little reference to this body of research.

basic contextual facts about labour market trends in developed countries, before briefly laying out the argument of the book and the organisation of the manuscript.

1.2 Some basic facts about employment, wages, and inequality in developed countries

1.2.1 Employment

Technological advances in the wider sense are to a large degree responsible for the tremendous improvements in living standards, life expectancy, and other indicators around the world in the past decades (Ritchie, Roser, Mispy & Ortiz-Ospina 2018). Yet, the optimism that seems so compelling when looking at these indicators of human progress has often been accompanied by a profound pessimism and anxiety when it comes to the impact technological progress may have on the realm of work.

This tension is hardly new. John Maynard Keynes famously predicted in a 1930 essay that within a century, the problem of economic scarcity might be solved and progressive societies would instead face the issue of technological unemployment (Keynes 2010). While according to Keynes, this was to be welcomed, since it would enable people to work no more than 15 hours per week and devote the rest of their time to culture and leisure, others never shared his optimism (see, e.g., White 1931).

While Keynes was right about the spectacular increase in living standards, his predictions on working time missed the mark and technological unemployment has not materialised on anything like the scale he imagined (Freeman 2008). Nevertheless, the fear that technological change may result in mass unemployment has recently experienced a resurgence. An influential study by Frey & Osborne (2017) finds that 47 percent of all US employment is at high risk of automation in the near future. For less advanced countries, this figure is often substantially higher still (Frey, Osborne & Holmes 2016). Coupled with the deindustrialisation that advanced economies have already experienced (Kollmeyer 2009), it is therefore unsurprising that many are worried what future technological change may mean for the ability of ordinary people to earn a decent living through work.

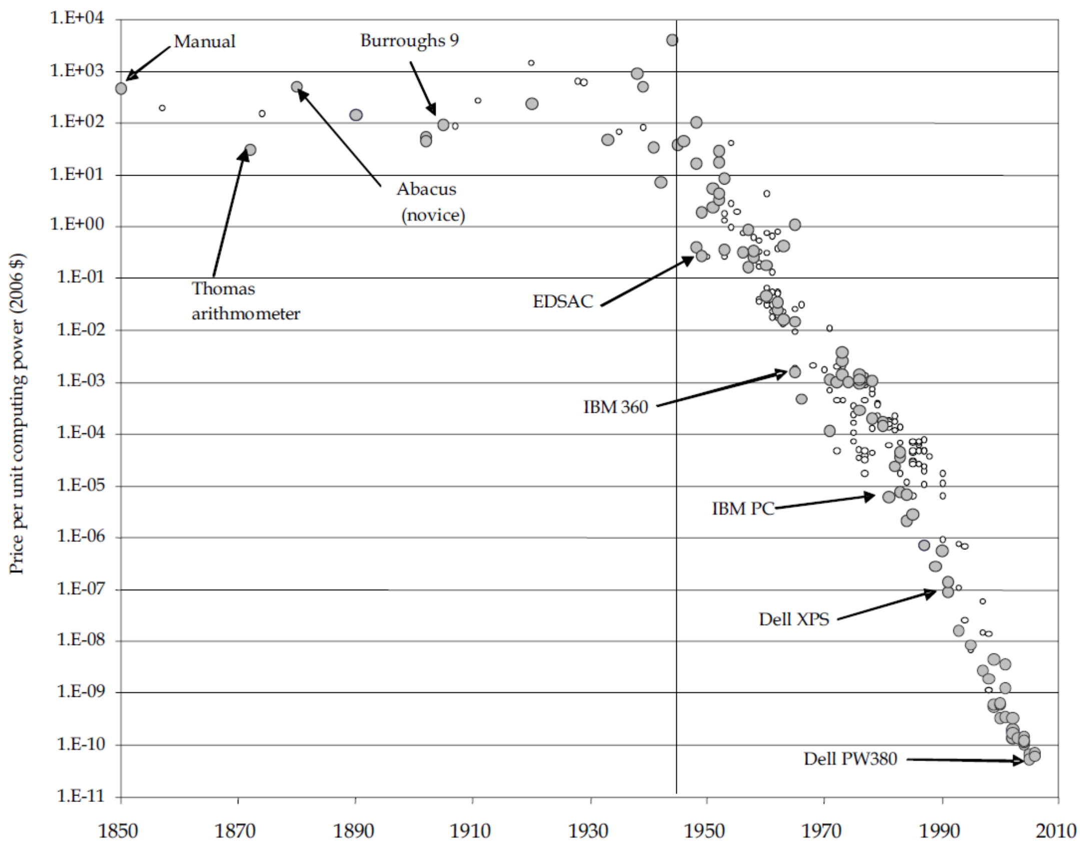

Figure 1.1: *The progress of computing measured in cost per computation per second deflated by the price index for GDP in 2006 prices. Source: Nordhaus (2007, p. 144).*

Just how momentous has technological progress been up to this point? The scale of technological progress can be expressed in numerous ways, but one particularly compelling perspective is William Nordhaus' (2007) study of productivity growth in computing. It shows that since 1945, the cost of computing has declined on average by an astonishing 45 percent annually. With no sign of this trend slowing down, as seen in figure 1.1, there are concerns that in fact, artificial intelligence may one day surpass human brains in general intelligence (Bostrom 2014). Even Nordhaus (2007, p. 158) surmises that perhaps, "computers will prove to be the ultimate outsourcer."

The exponential increase of computing power is a striking manifestation of technological change in the abstract, which has enabled many of the more tangible instances of technological progress in the workplace: the thoroughly digitalised economy of the present where approximately 40 percent of workers use digital technology at least half of the time (see table 5.3), but also ultimately the development of advanced machinery in the industrial sector. Thus, the continuous expansion of computing power is perhaps the most fundamental manifestation of technological change.

Reassuringly, until now, ongoing rapid technological progress has had little impact on aggregate employment. Unemployment rates have not generally increased in recent decades, and labour force participation has been stable, with small declines for men and small increases for women (OECD 2017). Work has not disappeared, but fundamental changes have nevertheless taken place. The employment structure in developed economies today looks very different from the distribution some 50 years ago; agricultural employment has almost disappeared and the relative share of manufacturing employment has declined substantially practically everywhere, while employment in the service sector has soared (DiPrete, De Graaf, Luijkx, Tåhlin & Blossfeld 1997, Kollmeyer 2009).[3]

Besides the sectoral changes, researchers are interested in which types of occupations grow and which disappear.[4] Across developed countries, there has been a near-universal increase in the employment shares of high-wage occupations in recent decades. There are, however, pronounced differences further down the occupational hierarchy. For polarisation, it is necessary that both low- and high-wage occupations expand their employment shares at the expense of medium-wage occupations.[5] For upgrading, it is necessary that high-wage occupations grow and low-wage occupations shrink, while medium-wage occupations may expand or decline. These stylised patterns are depicted in figure 1.2. In some cases of

[3]Economic sectors are classified based on standards such as the International Standard Industrial Classification of All Economic Activities (ISIC) or its European counterpart, NACE. Level 1 of ISIC/NACE comprises 21 different codes, but for simplicity a further aggregation into the three familiar sectors agriculture, industry, and services is common (United Nations 2008).

[4]Occupational classifications are more variable than the sectoral classifications; many national statistical offices employ their own scheme. The most commonly used standard is the International Standard Classification of Occupations (ISCO). The latest version is dated 2008, but the 1988 version (ISCO-88) is used throughout this thesis for data compatibility reasons. The concept of occupations is less straightforward than that of sectors. According to Elias (1997), occupations in ISCO are constructed around two key concepts: the job, defined as a set of tasks and duties, and the skills required for competent performance of the job. The ISCO classification thus contains both a task and a skill component.

[5]Throughout this thesis, polarisation generally refers to polarisation of employment in terms of occupational wages, unless otherwise noted.

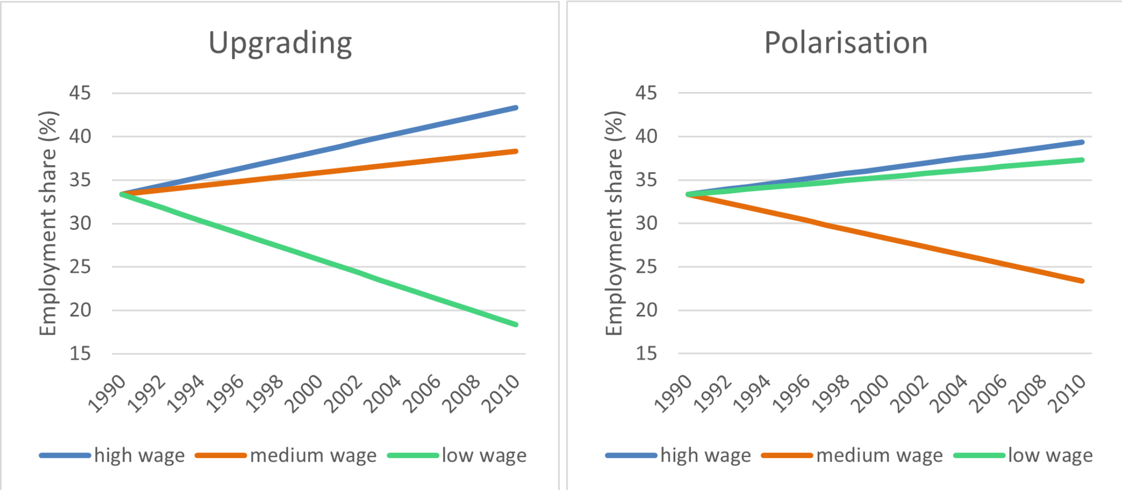

Figure 1.2: *Stylised patterns of employment change: upgrading and polarisation.*

upgrading, medium-wage occupations shrink faster than low-wage occupations. This pattern resembles polarisation, but since low-wage occupations also contract, I refer to it as polarised upgrading.

Figure 1.3 shows how these stylised patterns play out in the real world in a sample of OECD countries. This figure, taken from chapter 6 and based on data from the Luxembourg Income Study (LIS), shows polarisation in Germany, Luxembourg, The Netherlands, and the United States.[6] The other countries overwhelmingly show upgrading, except Hungary where low-wage occupations appear to have expanded substantially – a pattern that likely can be attributed to poor data quality. In Finland we could furthermore speak of polarised upgrading, as medium-wage occupations declined marginally more than low-wage occupations; the dominant feature is nevertheless the expansion of high-wage employment at the expense of the other terciles. This figure illustrates that there is great variation across countries regarding patterns of employment change. Other researchers using different data sources often find similarly diverse patterns, especially in European countries (Fernández-Macías 2012).

Explaining this variation is one of the core objectives of this thesis. In this section, I present some basic evidence about the various factors that have been

[6]The figure ranks 2-digit ISCO-88 occupations by their average wage in the respective country in the first period for which I have data and then classifies them so as to represent as close as possible to one third of employment.

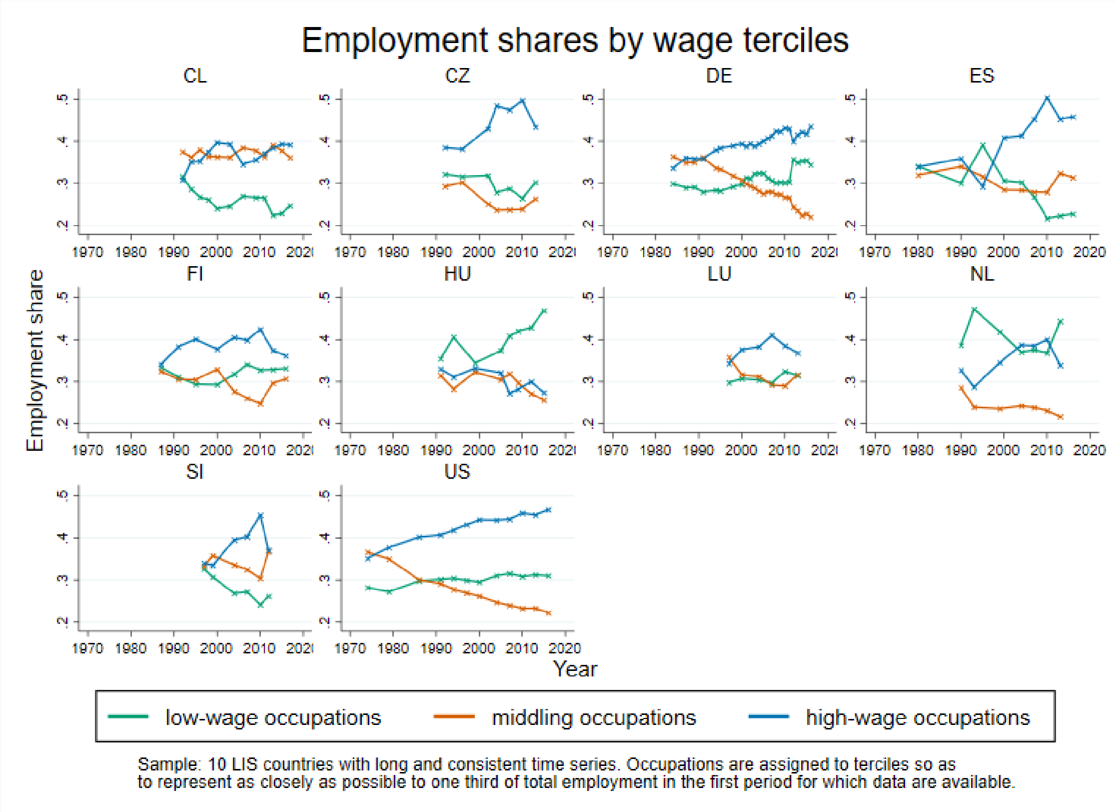

Figure 1.3: *Employment changes by occupational wage terciles.*

linked to employment and wage trends to allow the reader to situate my contribution in the wider literature.

Fuelling the fears of labour-replacing future technological change is the assessment that technological change has contributed to the large reallocations of employment between sectors and occupations in the recent past. The decline of low-skill, low-wage occupations or of semiskilled medium-wage jobs such as bookkeepers and the concomitant growth of service and professional employment have been attributed at least in part to technological change (Acemoglu & Autor 2011, Autor, Levy & Murnane 2003, Kollmeyer 2009, Mazzolari & Ragusa 2013).

Two related narratives dominate the discussion. Technological change has often been considered skill-biased (SBTC), complementing workers with higher educational attainment in a relationship that is consistent with the upgrading pattern of employment change (Goldin & Katz 2008). By contrast, other scholars have argued that information and communication technology primarily substitutes for labour in occupations that entail the performance of a large share of routine tasks (Autor et al. 2003). Based on the observation that these occupations are often medium wage, scholars advancing the RBTC hypothesis have predicted that the polarising pattern should prevail (Acemoglu & Autor 2011, Goos & Manning 2007, Goos, Manning & Salomons 2014).

Technology is of course not the only factor that may influence labour market trends. Globalisation and outsourcing have long been linked to the decline of the same occupations that have also been susceptible to technological change, in particular manufacturing occupations (Autor, Dorn & Hanson 2013). To complicate matters further, globalisation is itself at least partly a consequence of technological change which, by reducing the barriers associated with geographic distance, enabled the outsourcing of production in the first place (Goldfarb & Tucker 2019, Kollmeyer 2009). Figure 1.4 illustrates the pace of globalisation between 1980 and 2016 in a similar sample of countries as figure 1.3, operationalised by the sum of imports and exports as a share of GDP.[7] To assess globalisation, it is most instructive to look at the changes rather than levels of the trade-to-GDP ratio (OECD 2010). Here we see a pronounced increase almost everywhere from the 1990s until the Great Recession in 2009. After a recovery roughly to pre-crisis levels, trade as a percentage of GDP has been relatively stable in most countries, with small declines in Chile, Finland, and most notably the United States. It is therefore clear that the core period of globalisation coincides with the core period of employment polarisation.

Another crucial development that coincides partly with recent technological change is the massive expansion of female labour supply. Figure 1.5 shows the labour force participation rates of all women aged over 15 years from 1990 until 2016 in the same 10 countries as in figure 1.3. Except for Finland and the US, which already had high female LFP in the 1990s and where increasing enrolment in higher education may be counteracting the trend to increased LFP in younger cohorts (at

[7] Luxembourg is excluded here because its trade-to-GDP ratio far surpasses all other countries due to the reliance on financial services. Although it is no comprehensive measure of globalisation, the trade-to-GDP ratio is the most commonly used indicator of the importance of international transactions compared to domestic wealth creation (OECD 2010).

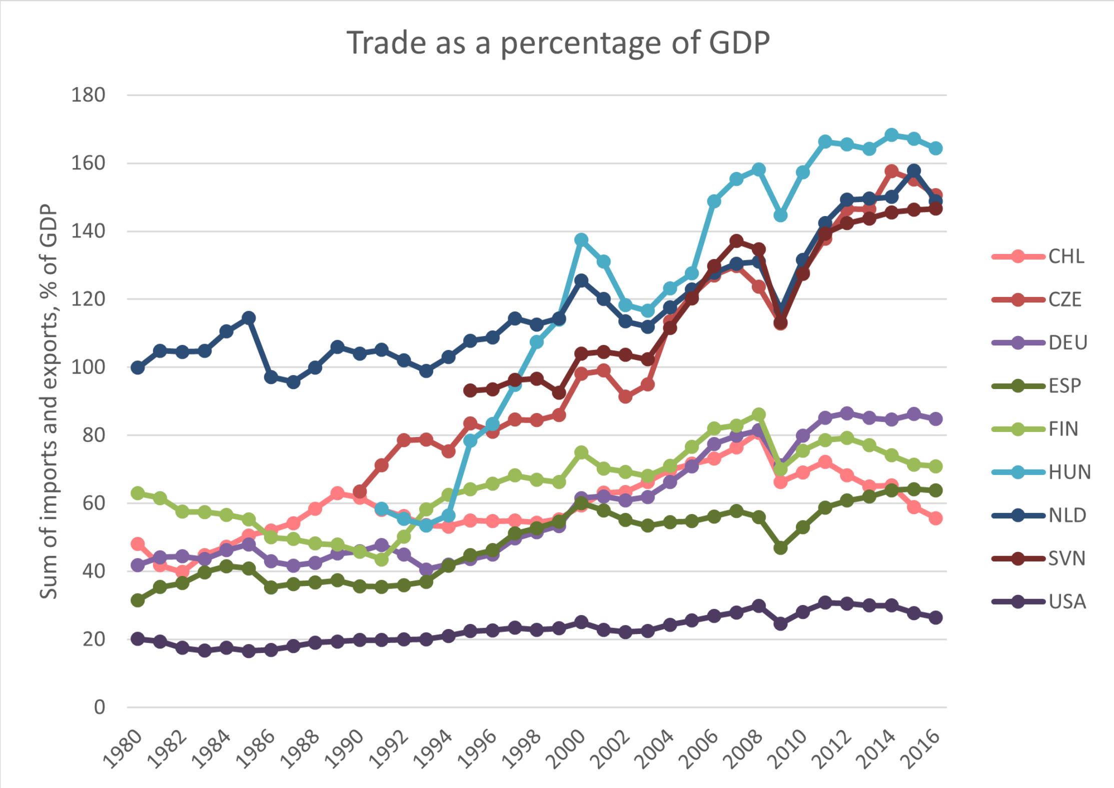

Figure 1.4: *Increase in the trade-to-GDP ratio, 1980 - 2016. Source: World Bank (2021).*

least in the US, see OECD 2017), all countries experienced an increase in female LFP that was often rapid, such as in Chile, The Netherlands, or Spain. Hungary, on the other hand, shows a pronounced drop after the fall of the Iron Curtain, only reaching the levels seen in the early 1990s again in the mid-2010s. The increase in female LFP itself has been attributed to diverse factors such as cultural changes (Fernández 2013, Goldin 2014) or an increase in the demand for female-oriented social skills in high-wage occupations (Cortes, Jaimovich & Siu 2018).

Linked to the integration of women into the labour market is the expansion of sectors and occupations that employ predominantly women in activities that previously were often provided by women in the home in what has been called the "care economy" (Dwyer 2013, Dwyer & Wright 2019). Services such as childcare and domestic help that are no longer provided in the household are in turn made affordable for purchase in the market by the extra income earned (Oesch 2013). Some authors argue that this process has contributed to the growth of (mostly low-skilled) service jobs that elsewhere has been attributed to technological change

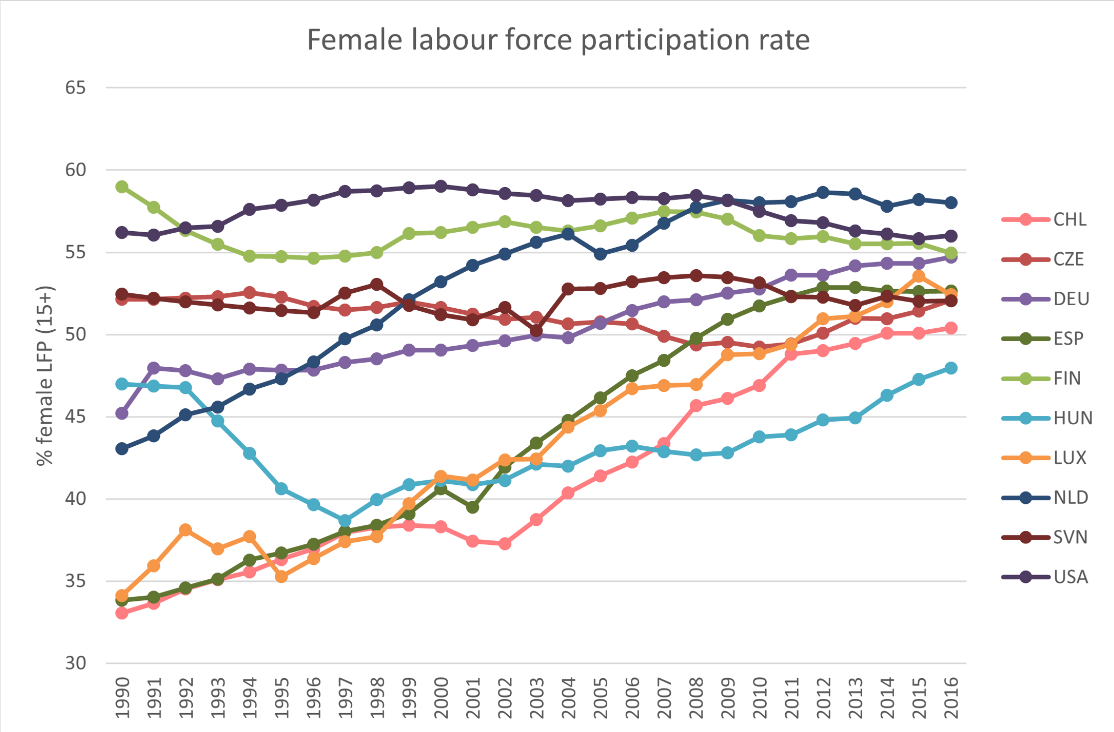

Figure 1.5: *Increase in female labour force participation, 1990 - 2016. Source: World Bank (2021).*

and consumption spillovers (Acemoglu & Autor 2011, Mazzolari & Ragusa 2013, Moretti 2012, Oesch 2013). Dwyer (2013) shows that this trend in the US also had a significant high-skilled component, making it one possible contributor to polarisation.

Lastly, labour market institutions may have played a role in the reorganisation of employment that has taken such different paths in different countries. Emerging employment polarisation in Germany, for example, has been linked to RBTC as well as to the liberalising labour market reforms of the early 2000s that are associated with employment growth and declining wages in the low-wage sector (Oesch 2015). Others contend that the decentralisation of wage bargaining and the ensuing period of wage restraint were to a large degree responsible for the employment boom in Germany in the 2000s (Dustmann, Fitzenberger, Schönberg & Spitz-Oener 2014). In any case, there is a strong case for an effect of institutions on the employment structure. This brief discussion of macro trends illustrates the close relationship between various labour market characteristics and their hypothesised impact on the employment structure. Moreover, the concurrence of different processes with similar putative effects makes it exceedingly difficult to disentangle them.

1.2.2 Wages and inequality

The factors just described are also linked to changes in wages and inequality. Overall inequalities in OECD countries have increased in recent decades according to various measures. The OECD reports that on average between 1985 and 2016, median incomes in 17 OECD countries have grown by less than 40 percent, compared to over 60 percent income growth for the top 10 percent; the incomes of the bottom 10 percent have grown slower still (OECD 2019). Dustmann et al. (2014) show that the same holds true for the growth of real wages in Germany between 1990 and 2008, while Machin (2016) finds a similar pattern in the UK. Thus, in many countries there has been real wage stagnation accompanied by rising wage inequality, which has resulted in stagnant living standards for the typical worker (Machin 2016). Figure 1.6, reproduced from Dustmann et al. (2014), illustrates the fanning out of the wage distribution in Germany that has been characteristic of developed countries since roughly the 1980s. This corresponds to an increase in the wage decile ratios that are widely used in the inequality literature (Checchi & García-Penalosa 2008).[8] The Gini coefficient for wages, another widely used measure of inequality, has likewise increased in Germany, the United States, and many other countries since the 1980s (Atkinson 2015, Checchi & García-Penalosa 2008).

In addition to overall increases in wage inequality, in some countries a polarisation of wages has been observed. Similar to employment polarisation, wage polarisation describes a situation in which wages at the top and the bottom of the wage distribution grow faster than in the middle. Most notably, this has been the case in the US, where simultaneous employment and wage polarisation have been documented, among others, by Acemoglu & Autor (2011), Antonczyk et al. (2018), Autor & Dorn (2013), Firpo, Fortin & Lemieux (2011). In most countries,

[8] Usually, the 50/10 ratio is used to measure inequality in the lower half of the wage distribution, while the 90/50 ratio captures inequality in the upper half.

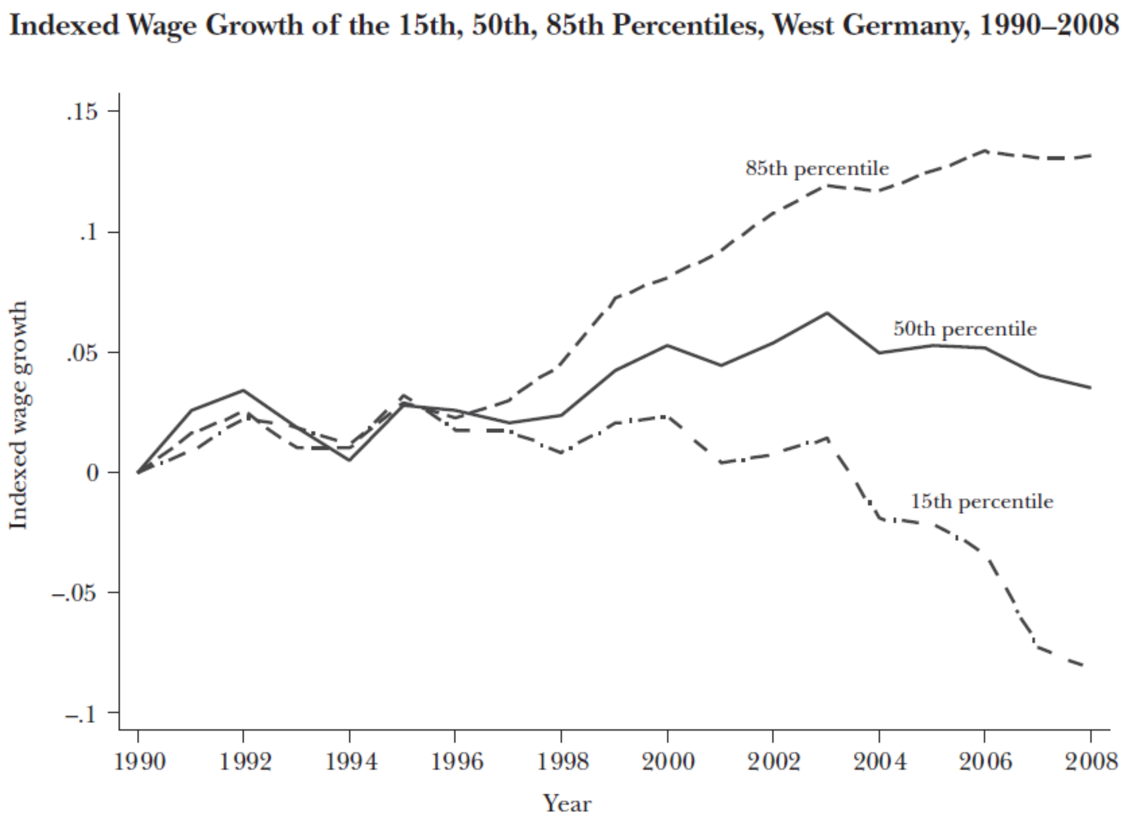

Figure 1.6: *Wage growth and increasing wage inequality in West Germany, 1990 - 2008. Source: Dustmann et al. (2014, p. 171).*

however, there is little evidence for wage polarisation, even if they have experienced employment polarisation (Naticchioni, Ragusa & Massari 2014).

Many of the factors discussed with regard to employment have also been linked to increases in wage inequality or wage polarisation. For example, technological change is considered one of the key drivers of inequality, whether it is polarising or not (Atkinson 2015). Even though Atkinson (2015) and others such as Mishel, Schmitt & Shierholz (2014) question this narrative, most scholars agree that technological change has contributed to increasing wage inequality by raising the demand for high-skilled workers faster than the education and training system could supply such workers, leading to an increase in the skill wage premium (Goldin & Katz 2008). Figure 1.7, reproduced from Acemoglu & Autor (2011), shows that the college wage premium in the US started a sustained increase in the early 1980s, when the supply of new college graduates began to slow down amid continued high demand.

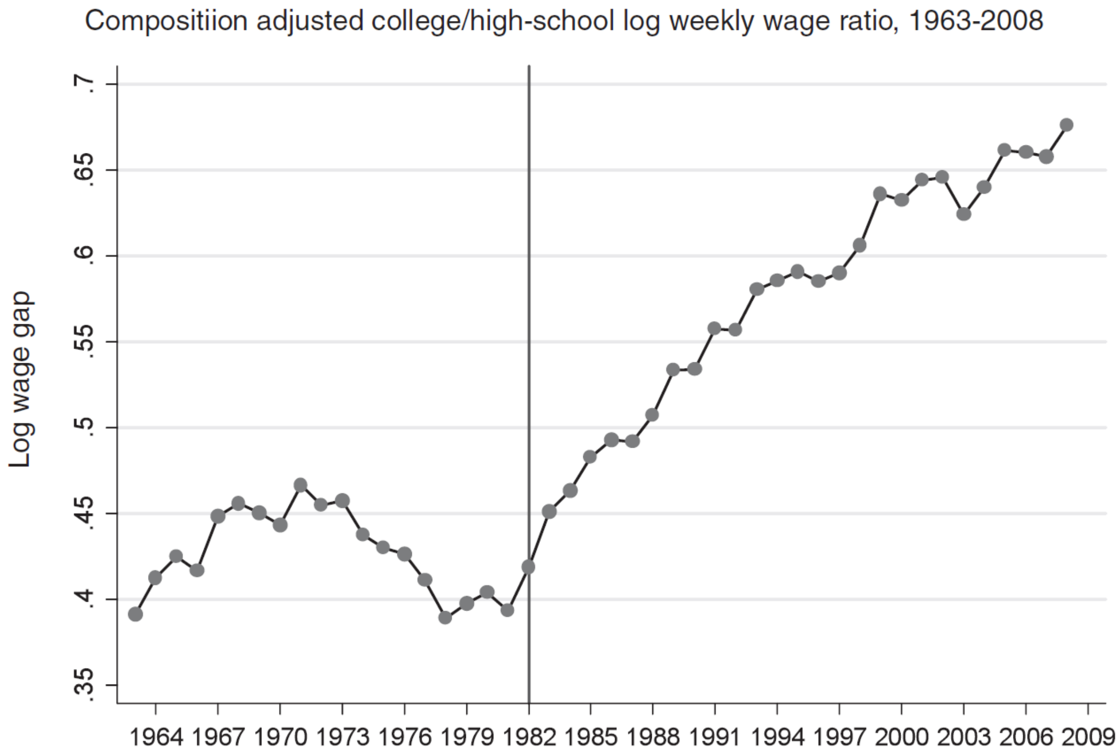

Figure 1.7: *The evolution of the college wage premium in the US. The line at the year 1982 marks the inflection point of the relative supply of college workers, which slowed down markedly after 1982. Source: Acemoglu & Autor (2011, p. 1052).*

Globalisation may also be linked to wage stagnation and inequality. At the job level, if the outsourcing of production is linked to employment losses, economic theory suggests that this should put pressure on wages in affected occupations. Blinder (2009) finds evidence for a wage penalty in the most offshorable jobs in the US in 2004 and Firpo et al. (2011) find that offshorability is linked to wage polarisation. At the country-level, trade integration, which I showed in figure 1.4 has increased dramatically, should be associated with higher relative wages for skilled workers according to the projections of the Heckscher-Ohlin model (Förster & Tóth 2015). However, this is not necessarily borne out in cross country analyses (OECD 2011). Thus, while the wage effects of globalisation remain ever contentious, it is at least a plausible contributing factor.

Increasing female LFP may likewise have an impact on wages and inequality. When women began to enter the labour market in large numbers, they were generally less educated than men (Vincent-Lancrin 2008). Moreover, women disproportionately entered low-wage service jobs, which would tend to increase the distance of these occupations from the median wage, and thus lower-tail wage inquality. More recently, however, women are on average more educated than men in OECD countries (Vincent-Lancrin 2008), and growth in female employment has been more concentrated in high-wage occupations than male employment growth (Dwyer 2013, Oesch 2015). This could in turn reduce wage inequality in the upper half of the distribution as well as between the sexes (Blau & Kahn 2017).

Finally, a very extensive literature across various disciplines links the erosion of labour market institutions to wage inequality. This has been particularly well documented in the case of labour unions, whose membership has declined almost everywhere since their heyday in the 1970s and 1980s (Ebbinghaus & Visser 1999, 2000). Figure 1.8 displays this secular trend in my sample of countries. The decline is especially pronounced in the post-socialist economies throughout the 1990s, but it is visible everywhere except in Spain. Even in the corporatist welfare states with a strong union tradition such as Germany, union density collapsed by almost 50 percent (Lehndorff, Dribbusch & Schulten 2018). By depriving workers of a unified representation in wage negotiations, union decline has been linked to increasing inequality in a number of studies (e.g. Card, Lemieux & Riddell 2004).

This overview illustrates that there are a multitude of factors that could play a role in shaping the broad labour market trends with regard to employment, wages, and inequality. The challenge for researchers is thus to strike a balance between comprehensiveness and parsimony. In this thesis, I focus on technological change and LMI as the two key drivers of labour market trends. This does not mean that I discount the importance of globalisation, increasing female LFP and sectoral changes; I discuss and account for these factors as appropriate.

However, technology and institutions emerge in the literature as the two most consistently important drivers of labour market change. Moreover, academic as well as political debates in recent years have pivoted towards technology and institutions at the expense of other factors like globalisation. In the case of technological change, this presumably reflects the profound importance of understanding the influence

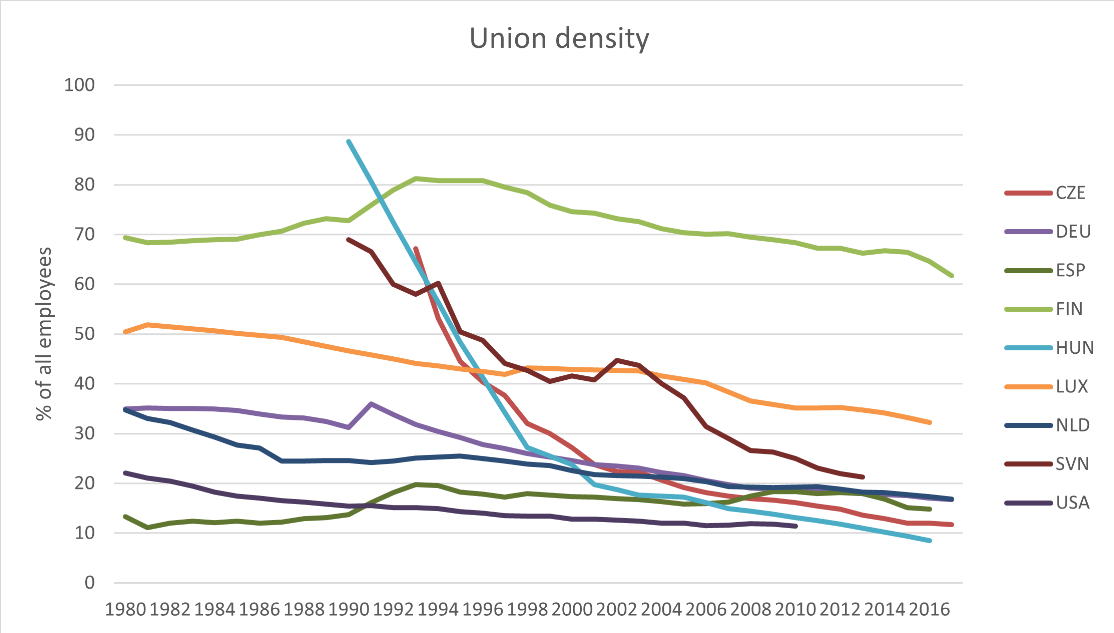

Figure 1.8: *Union density, 1980 - 2016. Source: Armingeon, Engler & Leemann (2020).*

of technology not only on the recent past but also on the future. Labour market institutions, on the other hand, offer the most direct lever for policy to affect desired outcomes and counteract undesired ones, should for example the pessimistic scenario regarding future technological change come to pass. For these reasons, focusing on the effect of technology and institutions on employment, wages, and inequality, is a sensible choice. With the scene thus set, in the remainder of this chapter I provide a preview of the argument and structure of this thesis.

1.3 Argument of the book

The previous section illustrated the complexity of the field that my research questions investigate. Numerous explanations exist for the phenomena of interest, but they have too often been treated as competing rather than complementary. Accordingly, the underlying argument of this book is that simplistic models in both economics and sociology have stood in the way of better understanding the labour market effects of technological change and LMI. My research questions are all linked to this core argument which pervades the thesis.

Firstly, to investigate the subsequent propositions, suitable measures of occupational tasks are required. Therefore, theoretically informed measures of occupational task content can address the first shortcoming of the existing literature, which is its reliance on outdated and unsuitable measures of occupational tasks.

Secondly, properly understood and operationalised, technological change is routine-biased as predicted by economic theories. There is also a skill-biased component that is related to, but distinct from the routine component. This element of my argument deals with a second simplistic assumption that prevails in the literature: SBTC and RBTC are often seen as alternatives rather than as complements.

Third, occupational wage hierarchies determine if RBTC plays out in an upgrading or polarising manner. The vast majority of existing research assumes that routine occupations are in the middle of the wage distribution (Goos et al. 2014). Some scholars argue that they are near the bottom (Fernández-Macías & Hurley 2017), but neither group investigates whether country differences in the occupational hierarchy could be behind the diverse patterns of employment change where RBTC should be at work. The argument that wage hierarchies matter, challenges the simplistic assumption that RBTC should everywhere lead to polarisation.

Finally, differences in occupational wage hierarchies themselves are partly influenced by institutional and technological factors. In particular, how well paid routine-intensive manufacturing occupations are differs systematically between polarising and upgrading countries. This contradicts the widespread notion that occupational hierarchies are invariant between countries (Hout & DiPrete 2006).

The empirical part of this thesis follows this structure and provides support for the individual elements of the overall argument. Thus, recent labour market trends are the result of both technological change and institutional factors; however, both economists and sociologists have blind spots when it comes to providing a unified theory. This thesis brings together the different literatures and based on this novel perspective develops important new insights.

1.4 Organisation of the manuscript

The thesis opens with an extensive discussion of the existing literature on the labour market effects of technological change and labour market institutions in chapter 2.2. The discussion aims to show on one hand how scholarly thinking about technological change has evolved in economics and sociology, and on the other hand to detail the remaining limitations which this thesis seeks to address. The first part of the chapter examines the evolution of the routine-biased technological change literature in labour economics and its origins in the skill-biased technological change framework. The first subsection shows the enduring importance of the argument that computer technology complements skilled workers, increasing the demand for and wages of such workers, and the refinements of this theory with the shift in emphasis from skills to job tasks. With this approach, scholars aimed to capture more directly what technology really does. The next subsection discusses the implications of RBTC for employment and wages, namely that jobs consisting of repetitive, codifiable tasks have been at the greatest risk of replacement by technology and real wage declines. Since the development of the SBTC and subsequently the RBTC model is essentially an American story, the first two subsections focus on research from the United States to outline the contours of the theories.

Following this, I trace the global proliferation of the RBTC paradigm and discuss the longitudinal and cross-country evidence outside the US. Since the evidence for pervasive polarisation outside the US is inconclusive, I also examine a growing body of sociological evidence that questions the predictions of RBTC theory. In this body of literature, a certain scepticism regarding the primacy of technology in explaining structural changes and increasing inequalities becomes apparent. For example, other structural changes such as rising female labour force participation and population ageing and the attendant rise in care work receive greater attention from sociologists. The sociological literature thus provides a helpful counterpoint to the sometimes excessive reliance on technological factors in the economics literature. In a fourth subsection, I moreover discuss the potential of globalisation to serve as an alternative explanation for the trends attributed to RBTC.

In addition to the economics literature and its sociological critiques, chapter 2.3 deals in detail with the role of labour market institutions in explaining recent labour market trends. In the form of power resource theory (Korpi 1983), the institutional literature offers what comes closest to a competing paradigm to RBTC. My analysis shows that while there is a rich body of research on the effect of unions and other LMI on outcomes such as unemployment and inequality, the relationship with relative occupational employment is surprisingly understudied. Furthermore, differential effects on specific occupational groups have received little attention. I therefore discuss some possible mechanisms how LMI may affect employment growth in specific occupational groups. Then, I go on to analyse what existing evidence tells us about the impact of LMI on the wages of groups such as manufacturing workers. Following this, chapter 2.4 briefly points out some further key limitations in the theoretical literature on both RBTC and LMI.

The theoretical chapter is followed by a brief chapter that lays out succinctly how this thesis advances existing theories. Based on the discussion in chapter 2, it sets out why there is a need for better measures of occupational tasks, why it is important to critically verify the assumptions that underpin the RBTC model, and why it is crucial to integrate technological and institutional arguments to explain differences in occupational wage hierarchies. The chapter sketches out the main thrust of the argument in each of the empirical chapters and shows how the chapters build on one another. Following this, chapter 3.4 discusses some common features of the analytical strategy throughout the individual chapters in this thesis. I describe my statistical approach and my rationale for choosing 2-digit ISCO-88 occupations as the basic level of analysis throughout this thesis, as well as my approach to sample selection and the associated challenges. In chapter 4, I introduce the main data sources for the thesis. I discuss both the advantages and disadvantages of the European Working Conditions Survey (EWCS) and the LIS and argue why they are best suited for my analyses. Chapters 3 and 4 therefore provide the unified methodological framework that ensures the overall coherence of the thesis.

The empirical part of this thesis addresses some important gaps in the literature on the labour market effects of technological change. Three main contributions have been identified above and thus structure the empirical part of the thesis. In chapter 5, I engage with the question how to best measure occupational task content. This is a crucial issue for the technological change literature because later findings are determined by the way routine occupations are defined and measured. Therefore, without a solid methodological basis, the whole edifice of the routine-biased technological change literature stands on shaky foundations. I propose two new measures of routine-task intensity and task complexity and compare them to prominent measures in the literature. The chapter argues that the new measures are more precise than the ones hitherto dominant in the literature. Substantively, it shows that there is significant variation between countries and over time that existing measures fail to capture.

In chapter 6, I show that despite the routine-biased nature of technological change, employment polarisation has not been as commonplace as claimed in much of the economics literature. I argue that the position of routine occupations in the wage hierarchy is key to understanding patterns of employment change in response to technological change. To this end, I propose a refined model of RBTC that accounts for country differences in the relative position of routine occupations. The chapter then introduces the concept of routine-wage curves, depicted in figure 1.9, to show that only where routine occupations cluster around the middle of the wage distribution are we likely to see polarisation. Where routine occupations are concentrated near the bottom of the wage hierarchy, upgrading occupational change is the norm. Based on research on the US, the former has been widely assumed, but it does not hold true in numerous European countries. I further show that routine-wage curves have been relatively stable over time, with richer countries more likely to exhibit a polarising pattern. This chapter therefore provides a straightforward explanation for the observed diversity of changes in the employment structure that can be reconciled with the core tenet of RBTC theory, that technology substitutes for workers in occupations that comprise a large share of routine tasks.

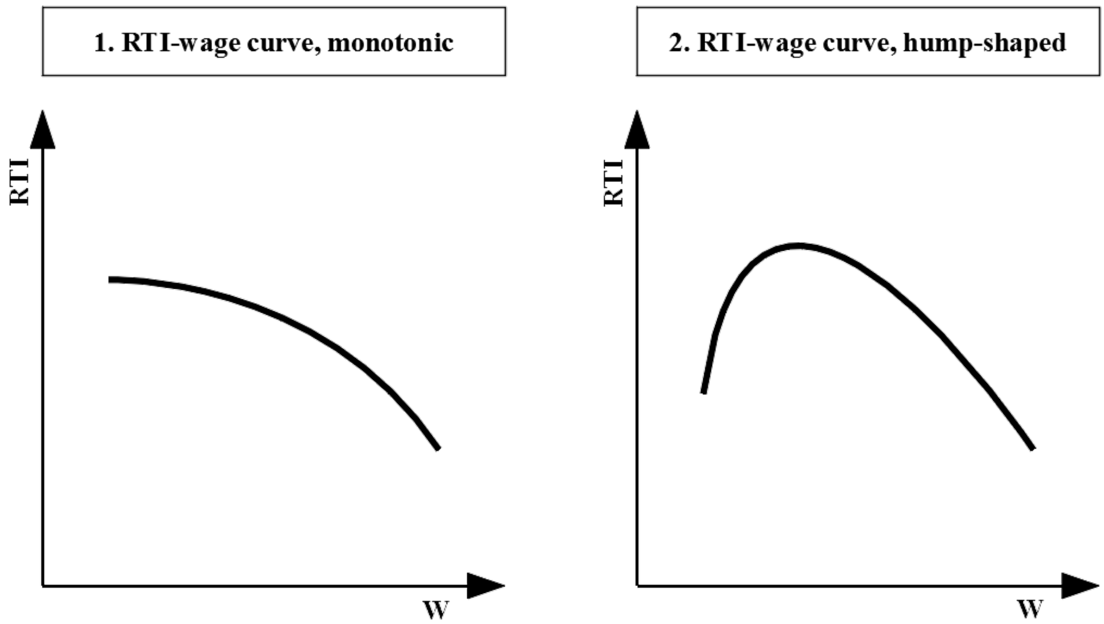

Figure 1.9: *This figure shows the stylised shapes of the routine-wage curves that are analysed in chapter 6.*

Chapter 7 takes the findings of the previous chapter and asks, why are there differences in the relative position of occupations with similar routine intensity? This appears to contradict well-established sociological findings regarding the stability of wage hierarchies across countries and over time (Hout & DiPrete 2006, Treiman 1977). Against this backdrop, the chapter investigates the role of industrial robots and labour market institutions in shaping occupational hierarchies, as expressed in the wage premium for routine manufacturing occupations compared to medium- to high-routine non-manufacturing occupations. I pursue three main arguments: firstly, that robotisation increases the relative productivity and hence wages of routine manufacturing workers, secondly, that unionisation (employment protection) increases (reduces) the relative bargaining power of routine manufacturing workers and hence the manufacturing wage premium, and thirdly, that LMI may moderate the effect of robotisation on relative wages. While I do not find evidence for the productivity hypothesis, the bargaining power hypothesis is partly supported, as stricter employment protection is associated with a lower manufacturing wage premium, and is furthermore linked to a stronger negative relationship between robot density and the manufacturing wage premium.

The conclusion in chapter 8 first summarises the empirical findings of the thesis and reiterates its contributions to the literature. I then go on to discuss the remaining limitations, focusing on the trade-off between comprehensiveness and parsimony in explaining complex phenomena, the drawbacks of using disparate data sources, and the potential to more fully use the task measures developed in chapter 5. Finally, I close by addressing the debate about the future of technological change and a potential "end of work" in light of my findings. I argue that fears of technological unemployment are likely overstated. While RBTC more broadly and robotisation in particular have led to a reallocation of employment, they do not so far appear to have significantly reduced aggregate employment, with no indication of this changing soon. Therefore, it would be more prudent to focus on improving the quality of existing jobs.

2

What moves labour markets?

Contents

2.1 Introduction

The introductory chapter outlined some key elements of the technological change debate, highlighting different emphases that have emerged regarding relevant worker or task characteristics, outcomes of interest, empirical mechanisms and alternative explanations. This chapter provides a more in-depth analysis and systematisation of this sprawling area of research.

It commences with an exposition in section 2.2 of the main theoretical framework that underpins this thesis, the routine-biased technological change hypothesis developed by labour economists. I trace its evolution out of the skill-biased technological change framework in the United States, its implications for employment and wages there, and its proliferation across disciplinary and geographical boundaries.

In this context, I also examine the limitations of the RBTC hypothesis, engaging above all with the various sociological critiques, and arguments that propose globalisation as an alternative explanation for the trends attributed to RBTC.

In section 2.3, I engage with the main alternative paradigm for analysing labour market outcomes, the power resource framework. I highlight existing research (or the lack thereof) on how unions and employment protection legislation (EPL) shape relative employment in routine occupations and incentives for technology adoption, as well as findings on potential distributional effects of these institutions. I furthermore discuss how this research relates to RBTC. It is from the tension between insights from economics and sociology that my hypotheses and research questions are born.

Finally, section 2.4 deals with the unaddressed limitations of both the RBTC and institutionalist frameworks.

2.2 Routine-biased technological change

2.2.1 Evolution

This section traces the evolution of the theoretical thinking about technological change and labour markets in economics. I focus on what I view as the paradigmatic models at various points in time to provide a narrative of how scholars of technological change have incorporated new evidence and updated their assumptions to arrive at the model of RBTC that is currently dominant in labour economics. I then follow up with a more detailed discussion of empirical evidence in the section on implications. Furthermore, since the development of the SBTC and subsequently the RBTC model is essentially an American story, these sections focus on research from the United States to outline the contours of the theories, whereas the section on proliferation tracks their application globally.

Factor-neutral technological change

Earlier theories of technological change had often treated it as factor-neutral (see, e.g., Solow 1957). Atkinson & Stiglitz's (1969) seminal article was one of the first papers to challenge the orthodoxy of the time that technological progress simply shifts the production function outwards. Instead, they posited that technological change may be localised, raising productivity only in certain technological applications.[1] Only slowly, this approach gained more widespread acceptance.

Skill-biased technological change

Against this backdrop, the understanding that technological change in the post-industrial era generally complements high-skilled workers was a major theoretical advance.[2] Investigating the surge in wage inequality in the US throughout the 1980s, economists in the early 1990s found that a substantial part of the increase could be explained by the growing relative demand within detailed sectors for more educated, skilled workers (Katz & Murphy 1992, Mincer 1991). Furthermore, wage growth was fastest in occupations with higher levels of computer use (Katz & Murphy 1992, Krueger 1991). The skill-biased technological change hypothesis was born and generated substantial research interest in the following years (see, e.g., Autor, Katz & Krueger 1998, Berman, Bound & Griliches 1994, Berman, Bound & Machin 1998, Bound & Johnson 1992).

The proponents of SBTC argued that technology increases the relative productivity of high-skilled workers. This increase in productivity then increases the relative demand for skilled labour if skilled and unskilled labour are imperfect substitutes (Acemoglu & Autor 2011, Katz & Murphy 1992, Violante 2008).[3] However, the mechanism by which technology affects productivity and thus demand is often not well specified. For example, Katz & Murphy (1992, p. 36) merely mention in passing

[1] Acemoglu, Autor, Dorn, Hanson & Price (2014) revisit their article and analyse the theoretical developments it has spurred, for example endogenous growth theory (Romer 1994) and the directed technological change literature (Acemoglu 1998). Yet, Levy & Murnane (1992) in their review of trends and explanations of earnings inequality still refer to "nonneutral technological change" as if to an anomaly, illustrating the persistence of the conventional view of technological change.

[2] As numerous authors point out, historically technological change has often been biased towards low-skilled workers, as during the industrial revolution (see, e.g., Acemoglu 2002, Autor 2015, Mokyr, Vickers & Ziebarth 2015).

[3] About this there is widespread agreement in the literature. Katz & Murphy (1992) estimate an elasticity of substitution between high-skill and low-skill workers of 1.4, Acemoglu & Autor (2012) put the elasticity between 1.6 and 1.8. The general consensus in the literature is that the elasticity of substitution ranges between 1.4 and 2.0 (Acemoglu & Autor 2012, Hutter & Weber 2017).

that technological changes "possibly associated with the computer revolution" may have raised the relative demand for educated workers.[4] It is further worth noting that Katz & Murphy (1992) only consider within-industry shifts in demand as reflecting SBTC, whereas other contributions acknowledge that technological change may also induce between-industry shifts (see, e.g., Goos et al. 2014). In any case, as demand for skilled workers increases, markets adjust by supplying more skilled workers or by raising the price of skilled labour (the skill premium), or a combination of both (Violante 2008). Note that in this literature, skills are generally taken to be positively but imperfectly correlated with education, with the skill premium (in the case of the US, for having a college degree) serving as a proxy for technological change (Acemoglu 1999).[5]

This dynamic gave rise to labour market trends during the second half of the 20$^{\text{th}}$century that have been characterised as a "race between education and technology" (Goldin & Katz 2008).[6] The authors show that for much of the 20th century, demand for skilled workers in the US has increased in response to technological change that was, on average, skill-biased. Katz & Murphy (1992) and Levy & Murnane (1992, p. 1372) argue that there was "a steady increase in the demand for skilled workers relative to unskilled workers" in the US in the 1970s and 1980s. Berman et al. (1994) find evidence for SBTC in the US manufacturing sector in the 1980s, with the employment share of production workers dropping by 15 percent and an increase in the employment of nonproduction workers. Vivarelli (2014, p. 147) concludes that "evidence in favor of the skill-biased nature of new technologies is large, robust, and applicable across OECD countries, economic sectors, and types of innovation." Therefore, the skill-biased nature of employment

[4]Note also that SBTC not only increases the productivity of skilled workers. Instead, technological change leads to a general increase in productivity, but relatively more so for skilled workers. For this reason, the SBTC framework struggles to explain real wage declines of lower skill groups as was later pointed out by Acemoglu & Autor (2011).

[5]Holmes (2017) points to problems with using education as a proxy for skills, such as the undervaluation of skills acquired through on-the-job training. This is an additional, if not often stated, reason for preferring a task-based framework over a skill-based theory of technological change.

[6]This expression, in turn, was coined by Tinbergen (1974).

changes in recent decades and the associated upgrading of the employment structure is widely documented and uncontroversial.

One implication of the race between education and technology is that if technology is winning the race, the relative price of skilled labour is bound to increase, leading to higher wage inequality. Levy & Murnane (1992) and Katz & Murphy (1992) were among the first to document the increase of the skill premium for college educated workers in the 1980s and thereby rising overall inequality.[7] Both papers attribute the change largely to supply and demand shifts consistent with SBTC: secular growth in the demand for skilled workers and a slowdown in the growth of the average educational attainment of the labour force.

Bound & Johnson (1992) argue that technological change can explain the majority of not only the increase in the college wage premium, but also of the increasing wage premium for older workers and the reduction of the female-male wage differential during the 1980s. They specifically test the technological explanation vis-à-vis alternative hypotheses such as the decline of union power and a slowdown in the rate of growth of the college-educated population. They find that while all processes contributed to the observed patterns, technology was their primary cause. The data of Krusell, Ohanian, Ríos-Rull & Violante (2000) likewise show an increase of the skill premium from the 1980s onward, after a slight decline in the 1970s. They argue that this is the result of a capital-skill complementarity effect driving the skill premium up by 60 percent between 1963 and 1991, and a countervailing relative quantity effect (educational expansion) which reduced the skill premium by 40 percent. Thus, although their central argument is that technological change is capital-embodied, i.e., that capital and high-skill labour are complements in production, their predictions are very similar to those of the canonical model of Goldin & Katz (2008).

The SBTC model has been a workhorse model in economics for decades now. It provides a compelling explanation for the empirical pattern of increasing employment in skilled occupations and a concurrent rise in the skill premium that characterised

[7]Although their focus is on the skill premium, Levy & Murnane (1992) already observe the incipient hollowing out of the occupational structure.

much of the second half of the 20th century in the United States. Towards the end of the century, however, as technological change became increasingly synonymous with computerisation, the classical SBTC model required some crucial modifications which gave rise to the theory of routine-biased technological change.

Towards the routine-biased model of technological change

Like the SBTC model, the RBTC hypothesis was formulated based on the US experience in a seminal article by Autor et al. (2003) and refined further for example by Acemoglu & Autor (2011). In their article, Autor et al. (2003, p. 1280) set out to better understand "what it is that computers do". The core of their argument is worth quoting in full and stipulates

> "(1) that computer capital substitutes for workers in carrying out a limited and well-defined set of cognitive and manual activities, those that can be accomplished by following explicit rules (what we term "routine tasks"); and (2) that computer capital complements workers in carrying out problem-solving and complex communication activities ("nonroutine" tasks)."

Thus, what really matters is the nature of the tasks workers perform on their jobs, rather than their skill or education level: "By measuring the tasks performed in jobs rather than the educational credentials of workers performing those jobs, we believe our study supplies a missing conceptual and empirical link in the economic literature on technical change and skill demand", argue Autor et al. (2003, p. 1281). The RBTC model could hence also be called a model of task-based technological change. It is interesting to note that at this stage, the authors do not yet appear to view RBTC as a new paradigm, but rather as a refinement of the SBTC literature.

In the RBTC model, there are two types of labour input (routine and nonroutine) and workers supply a combination of both types of labour in line with their productivity endowments. The model predicts that computer adoption is faster in more routine-intensive industries or occupations, and that such industries and occupations will show a larger rise in nonroutine labour input and a larger decline in routine labour input (Autor et al. 2003, p. 1291). Importantly, they also specify a concrete causal mechanism by which computer technology affects skill demand: the declining price of computer capital which is well documented in the literature (see Nordhaus 2007). Empirically, Autor et al. (2003, p. 1312) find with regard to education that "task change is antecedent to educational upgrading, rather than merely a reflection of it", and can explain a large share – 60 to 90 percent – of the increase in relative demand for college-educated workers in the 1980s and 1990s. This paper thus does not yet consider the implications for wages (or employment in terms of wages, which is what studies of employment polarisation look at), but it initiated the shift in emphasis from worker characteristics to job tasks and thereby set the stage for a hugely influential literature.

Acemoglu & Autor (2011) have refined the RBTC hypothesis further and provided perhaps the most comprehensive task-based model of routine-biased technological change. Moreover, their model explicitly links RBTC to employment and wage polarisation, which had proved problematic for the SBTC literature. The model features a clear distinction between skills and tasks, hence "workers of a given skill level can potentially perform a variety of tasks and, moreover, can change the set of tasks that they perform in response to changes in supplies or technology" (Acemoglu & Autor 2011, p. 1119). To enable the study of polarisation, this model furthermore incorporates three different groups of workers rather than two in the SBTC and Autor et al. (2003) models. Given its centrality for the literature, it makes sense to write out some key elements of the formal model. Thus, the production function for the unique final good is given by

$$Y = exp\left[\int_0^1 ln\ y(i)\ \partial i\right] , \tag{2.1}$$

with $y(i)$ the "service" or production level of task i. Markets are assumed to be competitive and there are three types of workers, L, M, and H, who supply low-skilled, medium-skilled, and high-skilled labour. Each task can be produced by all types of workers and capital, and thus has the production function

$$y(i) = A_L \alpha_L(i) l(i) + A_M \alpha_M(i) m(i) + A_H \alpha_H(i) h(i) + A_K \alpha_K(i) k(i). \tag{2.2}$$

The A terms represent factor-augmenting technology, the α terms are task productivity schedules. For example, $\alpha_L(i)$ is the productivity of low-skill workers in task i, and $l(i)$ is the number of low-skill workers allocated to task i; the other terms are defined analogously. While tasks can be performed by any type of worker, the comparative advantage of skill groups differs across tasks.

These basic equations illustrate some of the key characteristics that distinguish this model from the SBTC model as well as the earlier Autor et al. (2003) model. First of all, the final output is the product of a range of tasks rather than a combination of discrete labour and capital inputs as in Autor et al. (2003). Related to this is the specification of a task production function with three types of workers who can perform all types of tasks, but with differences in comparative advantage of skill groups across tasks, as captured by the α terms. The assumption regarding comparative advantage is simply that $\alpha_L(i)/\alpha_M(i)$ and $\alpha_M(i)/\alpha_L(i)$ are continuously differentiable and continuously decreasing. This implies that "there will exist some I_L and I_H such that all tasks $i < I_L$ will be performed by low skill workers and all tasks $i > I_H$ will be performed by high skill workers. Intermediate tasks will be performed by medium skilled workers." (Acemoglu & Autor 2011, p. 1122).

With this and some further assumptions, the model can account for the declining employment shares and relative wages of medium-wage routine workers compared to low-wage nonroutine workers that the SBTC and Autor et al. (2003) models could not explain. As Acemoglu & Autor (2011, p. 1137) state, "there is a central contrast between our framework and the canonical [SBTC] model: any improvement in technology in the canonical model raises the wages of all workers, whereas in our task-based framework an increase in A_H (high skill biased technical change), for example, can reduce the wages of medium skilled workers because it erodes their comparative advantage and displaces them from (some of) the tasks that they were previously performing." While these excerpts by no means fully characterise the model that is set out in Acemoglu & Autor (2011), they illustrate the most important features that are the reason why this model is widely seen as a reference point in labour economics.

This more flexible framework can account for crucial empirical observations such as the decline in medium-wage employment and increase in low-wage employment (polarisation), and concomitant real wage declines for medium-wage workers, that have characterised the US labour market in recent decades. Indeed, the predictions derived from this model correspond well to the employment and wage trends in the US in the 1990s and 2000s, in particular the erosion of relative employment and relative wages in medium-wage occupations, and it has since become the predominant way of thinking about technological and occupational change in developed countries in economics.

To recapitulate, as the RBTC hypothesis evolved out of the SBTC model, both share a similar logical structure. Technological change itself is often taken to be spurred by the decline of the real price of (computer) capital, which can be assumed to be comparable across developed countries (Autor et al. 2003, Koeniger et al. 2007, Spitz-Oener 2006). Often, this technological change is assumed to be exogenous, but it can also be modelled as an endogenous process (Acemoglu & Autor 2011).

The key difference is implied in the names of the theories: in the case of SBTC, it is worker characteristics that matter. Higher skills, usually operationalised with educational attainment, are said to be complemented by new technologies, and occupations which require more skilled workers are expected to expand. If educational expansion does not keep pace in the race with technology, wage premia increase, giving rise to growing wage inequality (Goldin & Katz 2008). In the RBTC model, the focus is on occupational tasks rather than worker characteristics. In the reasoning of Autor et al. (2003), computer capital substitutes for workers in "routine" activities that can be accomplished by following explicit rules and complements workers in carrying out “non-routine” tasks such as problem-solving and complex communications. Thus, employment is expected to contract in occupations where workers perform mostly routine tasks. Employment should grow in non-routine cognitive occupations – essentially high-skilled occupations – where the aforementioned complementarities lead to increased demand, and in simpler non-routine occupations, either through the reallocation of former routine workers or through spillover effects from an increased demand for personal services by the growing group of high-skilled workers (Acemoglu & Autor 2011, Manning 2004, Mazzolari & Ragusa 2013, Moretti 2012).

2.2.2 Implications

Implications for employment

In this section, I discuss to which extent the predictions of RBTC theory are borne out in the empirical evidence in the United States. There are two main empirical strategies for investigating the argument that RBTC reduces routine employment. Most commonly, scholars have used measures of routine intensity to infer the impact of technological change. A less used alternative approach relies on direct measures of technological change, such as the value of ICT capital or, more recently, robot usage.

The implications of RBTC theory for changes in the employment structure follow from its assumptions on the position of high-routine workers in the wage hierarchy. While Autor et al. (2003) did not develop this aspect in their 2003 paper, Goos & Manning (2007) in a paper on the UK argued that routine occupations often require medium levels of skill and training and therefore cluster around the middle of the wage distribution. This notion was then picked up by scholars in the US. For example, Autor, Katz & Kearney (2008) show that routine tasks are most prevalent between the 20th and 60th skill percentiles; Autor & Dorn (2013) find a maximum of routine task intensity at approximately the 30th skill percentile.[8] Hence, with medium-wage occupations most susceptible to substitution, the RBTC hypothesis predicts employment polarisation and can rationalise declining wages for medium-wage workers (Acemoglu & Autor 2011). In terms of its empirical implications, this is one of the most consequential elements of the theory.

These theoretical arguments have been well substantiated empirically in the United States. In addition to the landmark 2003 study, David Autor has been involved in a number of follow-up articles with other authors on the United States (Acemoglu & Autor 2011, Autor & Dorn 2013, Autor et al. 2008). These studies,

[8]Skill (the percentile rank of an occupation's mean years of education) functions here as a proxy for the occupational wage.

using data from the Current Population Survey (CPS) with sample sizes reaching into the millions or census data, and measures of occupational tasks based on the Dictionary of Occupational Titles (DOT), invariably report a polarising employment structure. Autor & Dorn (2013) is a particularly interesting study, as it investigates patterns of employment change in 722 commuting zones, showing that even at a regional level, higher initial routine intensity correlates with greater adoption of computer technology and greater reallocation of workers from routine into service occupations (employment polarisation), as well as wage polarisation.

Numerous other economists have also contributed to the evidence regarding employment polarisation in the US, zooming in on different aspects of RBTC based on the Autor et al. (2003) framework. For example, Mazzolari & Ragusa (2013) identify increased demand from high-earners for personal services as a contributor to employment growth at the bottom of the wage distribution. Thus, they close one of the gaps in the original framework, which had clear predictions for employment trends at the top and in the middle, but not at the bottom of the wage distribution. Relatedly, Moretti (2012) finds a substantial multiplier effect of high-skill, high-tech jobs: according to his estimates, for every such job, 5 local service jobs are ultimately created, both in skilled and unskilled services.[9] Further insightful studies by Cortes (2016) and Cortes, Nekarda, Jaimovich & Siu (2020) track the occupational mobility and wage patterns of routine workers in a polarising labour market, thus zooming in on the individual-level ramifications of RBTC. They suggest that a large share of the decline in routine employment can be explained by reduced inflows from unemployment and labour force non-participation.

One study that questions not the existence of RBTC but the degree to which it is responsible for the observed polarisation, is Siegel & Bárány (2018). They argue that employment and wage polarisation in the US started as early as 1950, and thus long before computerisation could possibly have played a role. They consequently advocate for a greater role for structural economic change in explaining polarisation. However, as suggested in chapter 1, such structural change may itself partly reflect

[9] Manning (2004) arrives at a similar conclusion for the UK.

technological change. Cortes, Jaimovich & Siu (2017) likewise report employment polarisation between 1979 and 2014, but find that advances in technology as captured by the increase in the stock of ICT capital explain only a small portion of this trend. Unlike most of the empirical research on polarisation, this study uses a direct measure of technological change which may explain some of the difference in the findings. Using robot density as a direct measure of technological change, Acemoglu & Restrepo (2020*a*) also find a negative effect on relative routine employment. Overall, economists have produced a wealth of evidence that technological change has reduced routine employment in the US and that the employment structure has polarised, partly as a result of this.

Implications for wages

In parallel to the literature on the employment effects of RBTC, an empirical literature on the impact of RBTC on wages has developed. Indeed, employment polarisation is often seen as a mechanism through which technological change has contributed to wage polarisation and, thereby, greater wage inequality in the US.[10] For instance, Autor et al. (2008, p. 301) are predominantly interested in wage inequality between 1973 and 2005, but argue that "[t]he roughly parallel movement of earnings and employment growth in each decade suggests that demand forces have played a key role in shaping wage structure changes during the inequality surge of the 1980s and the polarisation that followed."[11] Thus, the changes in relative employment are seen as key to understanding the development of relative wages.

Other studies that find wage polarisation in the US include Cortes (2016) who estimates that the routine wage premium has fallen by 17 percent from 1976 to the mid-2000s compared to the wage premium for nonroutine manual occupations, while the premium for nonroutine cognitive occupations has risen by 25 percent in the same period. Interestingly, Cortes (2016) also estimates the individual-level

[10] While not every increase in inequality is polarising, wage polarisation almost always results in an increase in overall inequality. Even though wage polarisation by definition entails a reduction of wage inequality in the lower parts of the wage distribution, faster growth of top wages means that wage polarisation in the US has been associated with an increase in overall wage inequality.

[11] This study is also interesting because it analyses polarisation before the RBTC framework was fully fleshed out in Acemoglu & Autor (2011).

returns to switching from a routine to a nonroutine manual or a nonroutine cognitive occupation. Consistent with the polarisation narrative, he finds significantly faster wage growth for both groups of switchers compared to stayers in the long-run.

Similar strong evidence for RBTC-induced wage polarisation comes from Ross (2017), who demonstrates that occupational wages between 2004 and 2013 responded to increases in routine intensity with a statistically significant wage penalty, while increases in abstract task intensity carried a premium. Even Siegel & Bárány (2018), despite their finding that RBTC is not the primary explanation for employment polarisation, also document long-term wage polarisation since the 1950s. The predictions of the RBTC model are thus borne out with regard to both employment and wages in the United States.[12] In light of this body of evidence, it is unsurprising that RBTC is widely accepted as a framework for understanding the effect of technological progress on the American labour market. However, as the next section shows, attempts to apply this theory across the developed world have not always been so successful.

2.2.3 Proliferation

As Berman et al. (1998) point out, for SBTC to be a plausible explanation for employment and wage trends, it must be pervasive across advanced economies due to high levels of international communication and trade. Thus, it would be "hard to imagine major productive technological changes occurring in one country without rapid adoption by the same industries in countries at the same technological level" (Berman et al. 1998, p. 1247). The same is obviously true for RBTC. Hence, studies of RBTC and attendant polarisation have proliferated outside the narrow context of American labour economics. In this section, I outline the propagation of

[12] Among the few fundamental critics of the RBTC model within economics are Mishel, Bivens, Gould & Shierholz (2013), Mishel et al. (2014), and Mishel, Shierholz & Schmitt (2013). They point to various shortcomings of the RBTC literature, including a weak or missing link between occupational and wage trends. In addition, they argue that after 2000 there was no job polarisation to begin with, rendering RBTC unsuitable to account for recent developments, and that much of the increase in inequality happened in fact happened within detailed occupations (Mishel, Shierholz & Schmitt 2013). Lemieux (2006, 2008) is also critical of technological explanations for increasing inequality.

RBTC theory around the developed world and across disciplines, and the empirical challenges the theory has encountered in this process. For outside the US, the evidence for RBTC-induced polarisation of employment and wages is less clear-cut: sociologists have received claims of widespread polarisation through RBTC with scepticism, and while many studies in economics find employment polarisation in European countries, the link between employment and wage trends that appears so clear in the American data seems largely absent.

Country studies beyond the United States

A number of economists find employment and wage polarisation in Germany and the United Kingdom, the best-known examples being Goos & Manning (2007) for the UK and Spitz-Oener (2006) for Germany. These studies apply the approach of Autor et al. (2003) in different settings and find similar patterns with regard to polarisation. Importantly, Goos & Manning (2007) first introduced the notion that routine tasks which are susceptible to automation cluster around the middle of the wage distribution. They document a polarising employment structure and estimate that this can explain one-third to one-half of the increase in UK log wage differentials between 1975 and 1999. However, Salvatori (2018) looks at the longer period from 1979 until 2012 and finds that while there was job polarisation in the UK, it is not primarily explained by RBTC. He therefore concludes that the UK experience differs substantially from that documented in the US, and suggests that changes in the composition of the work force might be the dominant driver of polarisation.[13] In a related study, Cortes & Salvatori (2019) argue that increasing specialisation in non-routine occupations explains much of the increase in employment shares in these occupations in Britain.

Spitz-Oener (2006) finds that the employment structure in Germany has polarised between 1979 and 1999 based on a measure of occupational skill requirements. More importantly, she presents direct evidence that computerisation contributed to occupational task changes, offering a look inside the "black box" of technological

[13] Recall that compositional factors were found to be markedly less important than technological change in American studies such as Bound & Johnson (1992).

change and evidence for one of the mechanisms postulated by Autor et al. (2003). For the same period, Rendall & Weiss (2016) find that even though there was employment polarisation, the regions with the least routine employment had the highest rates of computer adoption, in apparent contradiction to the hypothesis that computers replace routine employment. They point out that the German apprenticeship system produces highly productive routine workers, making it costly to replace them with machinery. This varieties of capitalism-inspired explanation highlights the importance of institutions for employment outcomes.

Antonczyk et al. (2018), who compare wage inequality in Germany and the US between 1979 and 2004, find patterns that are consistent with a technology-driven polarisation of the employment structure. Yet, they caution that "the patterns in wage inequality in the two countries differ strongly, so that it is unlikely that technology effects alone ... can explain the empirical findings" (Antonczyk et al. 2018, p. 29). Dustmann, Ludsteck & Schönberg (2009), who like Antonczyk et al. (2018) use the Institute for Employment Research Employment Samples (IABS) which provide a 2 percent sample of social security records rather than the much smaller German Socioeconomic Panel (GSOEP) that is usually used, argue that the top half of the German wage structure did evolve in a way that is consistent with RBTC between 1975 and 2004. Their caveat, similar to Antonczyk et al. (2018), is that the pattern at the bottom differs from the US experience and may be better explained by institutional changes and other episodic events. This suggests that even though economists agree that there was polarisation in the UK and Germany, at least when it comes to employment, there are doubts to which extent this can be attributed to RBTC, and the findings with regard to wages are less conclusive still.

Since the German and British economies are close to the technological frontier, similar to the US, it is important to note that employment polarisation has also been documented in less advanced developed countries such as Spain (Anghel, De la Rica & Lacuesta 2014, Sebastian 2018) and Portugal (Fonseca, Lima & Pereira 2018). Anghel et al. (2014) replicate many of the findings of studies of the US or UK, including that changes in the composition of the labour force and employment reallocation between sectors cannot explain the observed pattern of polarisation between 1997 and 2012. By contrast, occupational routine intensity emerges as a powerful predictor of employment changes. Sebastian (2018) finds that the effect of polarisation between 1994 and 2014 on individual labour supply depends on degree status: routine workers with a degree tend to move up the occupational hierarchy, those without a degree, down. In Portugal, employment and wages polarised from the mid-1990s onwards in a manner that is consistent with RBTC (Fonseca et al. 2018). Interestingly, while they document a sharp decline in routine manual employment, Fonseca et al. (2018) find only a modest decline of routine cognitive employment, which they attribute to slower computer capital adoption. This of course contradicts those critics of the RBTC literature who argue that polarisation is essentially a phenomenon of liberal labour markets, while at the same time showing that country characteristics matter for how RBTC manifests itself.

Comparative studies in economics

In addition to these individual-country studies, a small number of comparative studies find a pervasive pattern of employment polarisation. Especially the article on 16 Western European countries between 1993 and 2010 by Goos et al. (2014) is widely cited in the economics literature and beyond. This paper has been so influential because it purports to show that the processes that have been identified in the US are also at work across Europe, resulting in pervasive job polarisation. Their model is similar to Acemoglu & Autor (2011) and constitutes the first comparative implementation of the RBTC framework proper. The paper finds sizeable between- and within-industry components of overall employment changes, more than two thirds of which can be explained by their model. This stands in contrast to some earlier models (e.g. Katz & Murphy 1992) which only consider within-industry changes as reflecting technological change. They also investigate the role of offshoring as an alternative explanation to RBTC, but find that RBTC is much more important for explaining employment polarisation. Goos et al. (2014) do not investigate the evolution of relative wages alongside employment shares, but nevertheless, this article has been among the most influential studies of RBTC and has inspired numerous follow-up analyses.

Michaels, Natraj & Van Reenen (2014) study 11 industrialised countries (US, Japan, and 9 European countries) from 1980 to 2004 and find that across countries, industries with faster ICT adoption increased their demand for highly educated workers relative to middle-educated workers in a way that is consistent with RBTC. An analysis of wage bill shares allows them to quantify both changes in hours worked and hourly wages. They find a polarising pattern in both, with a stronger adjustment in hourly wages (wage polarisation) in the US and stronger adjustments in hours worked (employment polarisation) in European countries. Another panel study of EU countries from 1995 – 2007 by Naticchioni et al. (2014) finds similar patterns at the industry level: changes in wage bill shares are consistent with polarisation, but they are almost fully explained by changes in hours worked. Thus, they find no evidence for wage polarisation operating through changes in relative hourly wages, either at the industry or at the individual level.

The OECD's (2015) flagship report on inequality sheds light on another facet of polarisation. As the report shows employment polarisation from the mid-1990s to 2010 in 11 out of 19 European countries, the authors point out that non-standard work has contributed significantly to this trend.[14] While employment losses in routine-intensive middling jobs have predominantly affected standard contracts, growth in abstract and non-routine manual work has often been in non-standard employment and at the extremes of the wage distribution.[15] Thus, the authors conclude, RBTC, while important, cannot be the full story behind polarisation (or the lack thereof).

A recent paper by Longmuir, Schröder & Targa (2020) is more in line with the findings of Goos et al. (2014). This study is particularly interesting in the context of

[14]Of the eight upgrading countries, four are Eastern European, while only one Eastern European country - Hungary - shows moderate polarisation.

[15]Related to this, Horemans, Marx & Nolan (2016) report that part-time employment, both voluntary and involuntary, increased in most European countries during the crisis from 2007 to 2013. However, whether part-time work is associated with polarisation depends on who disproportionately works part-time.

my thesis because it also uses LIS data. The authors find widespread polarisation: 30 out of 35 countries in their sample have experienced polarisation, including all developed countries except Slovakia. However, they investigate employment polarisation in terms of routine intensity rather than wages like the rest of this literature. Therefore, the results are not directly comparable. Like Mishel, Shierholz & Schmitt (2013) and Hunt & Nunn (2019), Longmuir et al. (2020) find no evidence for a close link between employment and earnings polarisation. Their paper thus provides important additional evidence from a data source that has hitherto rarely been exploited for this type of research.

These comparative studies, alongside the country-studies discussed earlier, thus paint a picture of widespread employment polarisation due to RBTC in developed countries. RBTC and employment polarisation are practically treated as equivalents: the possibility that RBTC could lead to non-polarising employment change is not seriously discussed. This follows from the view that routine occupations are predominantly medium-wage and the expectation that high-earners will increase their demand for low-wage personal services.

Whether wages follow the same pattern, as postulated by Acemoglu & Autor (2011) and verified by numerous studies for the US, is however highly questionable. The authors generally point to the different institutional setup that distinguishes European economies from the US and makes (wage) polarisation less likely (Antonczyk et al. 2018, Michaels et al. 2014, Naticchioni et al. 2014). To further complicate the picture, there seems to be somewhat of a disciplinary cleavage. So far, I have focused on the literature in economics; as the remainder of this section shows, numerous sociologists, especially in Europe, question the salience of RBTC as an explanatory framework for labour market change even where they have adopted elements of the biased technological change approach.

Sociologists on employment polarisation

In contrast to economists, sociologists have traditionally looked at social inequalities and changes in the employment structure through the lens of social class (see, e.g., Goldthorpe, Llewellyn & Payne 1987, Goldthorpe & McKnight 2004, Oesch 2006). Mostly concerned with social mobility rather than wage inequality, this literature generally shows a long-lasting process of class upgrading, which has however slowed down in recent years (see, e.g., Bukodi & Goldthorpe 2019, Goldthorpe 2013, for evidence for Britain).

Technology, while recognised as a driver of changes in the class structure, has been less central to this literature. However, some sociologists have increasingly adopted elements of the SBTC framework as a complementary perspective to the traditional class-based analyses (see, e.g., Fernández-Macías 2012, Oesch 2013, Oesch & Rodriguez Menes 2011). This includes a move to occupations as the unit of analysis and an explicit consideration of the factors that determine changes in the occupational structure, including technological progress. The concept of biased technological change has thus arrived in the sociological literature as well.

Sociologists who have adopted the occupation-based approach point out some important theoretical weaknesses of the models employed by labour economists and refine their predictions. For example, studies in economics rarely discuss how RBTC relates to other macro trends such as the emergence of the "care economy", rising female labour force participation, or immigration (Dwyer 2013, Eurofound 2017, Oesch 2015).[16] Yet, these developments affect the demand for and supply of labour largely independent of technology. Numerous sociological studies find a large effect of these trends - which are sometimes widely shared across countries, sometimes quite idiosyncratic - on changes in relative employment, especially towards the bottom of the occupational hierarchy. In this context, the RBTC model has been criticised for its sweeping predictions with little regard to country differences (Fernández-Macías 2012). In light of the cross-sectional variation regarding these factors, it appears unlikely that RBTC would everywhere lead to polarisation. Based on these arguments, several sociologists criticise what they consider an overemphasis on technology in some of the labour economics literature. Properly incorporating other

[16]One exception is Salvatori (2018) who finds a small role for immigration in explaining low-wage job growth in the UK.

relevant factors, the argument goes, polarisation either vanishes, or at least RBTC is dethroned as the main underlying cause of polarisation.

Importantly, the presence of employment polarisation in the US is not disputed by sociologists – in fact, the sociologists Wright & Dwyer (2003) were among the first to document emerging polarisation in the US in the 1990s. However, both in their 2003 paper and in later studies, Wright and Dwyer argue that RBTC is not the full story behind employment polarisation, pointing to the importance of the care economy in explaining employment growth in low-paying jobs (Dwyer 2013, Dwyer & Wright 2019).

Yet, studies claiming to show pervasive employment polarisation outside the US have generally been received with scepticism. Fernández-Macías (2012) criticises an earlier version of Goos et al.'s influential European study (Goos, Manning & Salomons 2009), and shows that its findings depend on a number of questionable methodological choices, only some of which are addressed in the final paper.[17] Thus, when Fernández-Macías (2012) looks at an almost identical sample and time period (EU-15 between 1995 – 2007), he finds job polarisation only in a handful of countries, while a majority show upgrading. Although Germany and the UK are among the countries exhibiting polarisation, this study suggests that polarisation is far from pervasive.

Following the same approach as Fernández-Macías (2012), Eurofound (2014, p. 31) finds that the overall pattern of employment change in Europe in 2011 – 2013 "was one of upgrading with some polarisation", with more upgrading changes for women and polarisation for men. Eurofound (2017), looking at the period from 2011 – 2016, finds more of the same diversity, and even downgrading in countries such as Hungary. Even though these studies only cover a relatively short period that is characterised by exceptional economic circumstances in the wake of the Great

[17] The three key differences Fernández-Macías (2012) lists relate to the definition of jobs, the job quality rankings, and the construction of job quality tiers. Not listed is the exclusion of agricultural and public sector-heavy occupations, even though Fernández-Macías (2012) uses the full working population in his analyses. The job quality (i.e., wage) rankings are country-specific in Goos et al.'s 2014 paper, and in some of the analyses jobs are defined as occupation-industry cells, but the problems of the uneven job quality tiers and excluded occupations remain. Thus, the 2012 critique for the most part still applies to the 2014 paper.

Recession, their findings cast further doubt on the narrative of straightforward polarisation.[18] Fernández-Macías & Hurley (2017) in an important paper go beyond documenting these differences and argue that routine occupations in fact cluster near the bottom of the wage distribution and that RBTC therefore should lead to the same upgrading employment change as SBTC. Where polarisation does occur, they therefore posit that it is not primarily due to technological factors. These studies do not reject the distinction between routine and non-routine jobs as such, but they caution against sweeping and overly deterministic claims regarding the employment effects of RBTC. Fernández-Macías & Hurley's 2017 article is particularly insightful since it is the first to question the equation of routine work with medium-wage jobs.

A number of case studies add to this evidence over a longer time period, even though some of their detailed findings differ. For example, using national labour force surveys, Oesch & Rodriguez Menes (2011) find polarisation in the UK and upgrading in Germany, Spain, and Switzerland in the 1990 – 2008 time frame. Oesch (2013), too, finds dominant upgrading in Denmark, Germany, Spain, and Switzerland in the 1990s and 2000s, and polarisation only in Britain. Oesch & Piccitto (2019) use the EU-LFS for the period from 1992 to 2015 and arrive at very similar results: polarisation in the UK and upgrading in Germany, Spain and Sweden. Finally, Murphy & Oesch (2018) investigate long-term (1970 – 2010) labour market trends in Ireland and Switzerland and likewise find no evidence for a simple story of upgrading morphing into polarisation.

However, while most sociologists agree that polarisation has not been as common as claimed by Goos et al. (2014), the example of Germany shows that there is still no consensus as to which countries have experienced it. Using the same data source (EU-LFS), Fernández-Macías (2012) finds polarisation between 1995 and 2007, and Oesch & Piccitto (2019) upgrading between 1992 and 2015. The different time

[18] In a detailed study of Ireland, Nolan & Voitchovsky (2016) find that the risk of job loss during the Great Recession was negatively related to monthly earnings, showing how idiosyncratic events can affect employment trends that may otherwise be shaped by technological change.

periods certainly play a role, as may other details of the operationalisation.[19] The evidence from these studies strongly suggests that numerous factors broadly related to the composition of the workforce, which have been largely ignored in the labour economics literature, have contributed to employment trends, which in European countries has often resulted in patterns other than polarisation. In addition to this, as I show below, sociologists have also made important contributions to the institutionalist literature which attributes employment and wage trends to labour market institutions rather than RBTC.

Sociologists on wage polarisation

Sociological studies of RBTC and polarisation mostly focus on the employment aspect. Wage polarisation is, at best, of secondary importance. This seems logical given that wage polarisation in the causal chain of RBTC theory is predicated on there being employment polarisation, and most sociologists dispute widespread employment polarisation in the first place. As Eurofound (2017, p. 33) explain, "other things being equal, a polarised labour demand would reduce wage inequality in the bottom half of the distribution (since the wages of those at the very bottom would increase with demand, relative to those in the middle), while increasing it at the top half." However, in the absence of polarised labour demand, there is no reason to expect wage polarisation.[20] The lack of a large sociological literature on wage polarisation, despite there being a substantial amount of research on the impact (or lack thereof) of RBTC on employment, indicates that this is not viewed as a major issue.

This does not mean, however, that sociologists have nothing to say about wage inequality in general, and the treatment of wages in the RBTC model in particular. For example, several sociologists criticise the absence of labour market institutions from the model, even though a large literature, including in other strands of

[19] For example, an interesting recent paper shows that in the US, the polarising pattern is stronger when agricultural workers are excluded (Cerina, Moro & Rendall 2021). Such details are not often discussed but may significantly influence results.

[20] Mishel & Bivens (2017) and Mishel, Shierholz & Schmitt (2013) argue with reference to the US that the development of an occupation's employment share is unrelated to changes in its relative wage and hence, even if there was job polarisation, wage polarisation would not follow.

economics, finds these to be important determinants of wage trends (Koeniger et al. 2007, Kristal & Cohen 2015, Parolin 2021). Parolin (2021, p. 924) states that "the relegation of institutions [in the RBTC literature], … is surprising given the rich history of literature on their role in shaping the earnings distribution within and between occupations", and Kristal & Cohen (2017, p. 187), following Goldin & Katz (2008), speak of a "race between institutions and technology" in which the former emerge victorious when it comes to explaining rising wage inequality in the US. The sociological literature therefore resembles the institutionalist literature when it comes to explaining changes in occupational wages, polarising or not.

Thus, extended beyond the United States, the empirical predictions of RBTC theory, widespread polarisation of employment and wages, are only partly borne out. Yet, it is in the nature of this sort of general theory that it should be universally applicable in countries at a similar level of economic development (Berman et al. 1998). Economic and sociological studies critical of the RBTC literature point to other factors such as the increased labour supply of educated women, low-skill immigration, and labour market institutions as alternative explanations for the observed diverse patterns of employment change and wage inequality.

To conclude this discussion, table 2.1 below shows an overview of studies of employment change, categorised by their scope (single-country study or comparative [case] study) and the patterns of employment change which they find.[21] The table shows that among the single-country studies, the findings lean heavily towards polarisation, while comparative studies more often report upgrading. This may reflect bias against the publication of single-country studies which question the dominant polarisation narrative. Furthermore, economists tend to find employment polarisation, while economic sociologists (and more institutionally minded economists) find more diverse, but predominantly upgrading patterns of employment change. This diversity is usually attributed to institutional and structural-demographic factors which either overpower the effect of RBTC or operate instead of it. Moreover, detailed methodological choices likely play a role where

[21] While many of these studies also investigate wage polarisation, not all of them do, hence the table focuses on employment.

studies find contradictory results for the same country, such as Fernández-Macías (2012) and Oesch & Piccitto (2019) in the case of Germany. As a result of these diverse findings, there remains considerable uncertainty how and to which extent RBTC has affected employment trends in developed countries, except in a few very clear cases such as the US.

Table 2.1: Overview of studies of employment change

	Single country studies	**Comparative studies**
Unambiguous polarisation	DE: Rendall & Weiss 2016 ES: Anghel et al. 2014 Sebastian 2018 PT: Fonseca et al. 2018 SE: Heyman 2016 UK: Goos & Manning 2007 Cortes & Salvatori 2019 US: Autor et al. 2008 Autor & Dorn 2013 Mazzolari & Ragusa 2013 Cortes 2016 Cortes et al. 2017 Siegel & Bárány 2018	Goos et al. 2014 Michaels et al. 2014
Dominant polarisation	DE: Spitz-Oener 2006 Dustmann et al. 2009 SE: Adermon & Gustavsson 2015 UK: Gallie 1991 Salvatori 2018 US: Acemoglu & Autor 2011	Naticchioni et al. 2014 OECD 2015 Antonczyk et al. 2018 Cirillo 2018 Mahutga, Curran & Roberts 2018 Jerbashian 2019 Longmuir et al. 2020
Dominant upgrading	CH: Balsmeier & Woerter 2019	Oesch & Rodriguez Menes 2011 Fernández-Macías 2012 Eurofound 2014, 2017 Fernández-Macías & Hurley 2017 Murphy & Oesch 2018
Unambiguous upgrading	SE: Tahlin 2019 US: Katz & Murphy 1992 Berman et al. 1994 Katz & Autor 1999	Berman et al. 1998 Oesch & Piccitto 2019

2.2.4 Globalisation and a changing nature of technological change

Before moving on to the literature on labour market institutions, a brief discussion of the primary alternative structural explanation for the decline of medium-wage employment - globalisation - is appropriate. The argument here is that the occupations that are susceptible to automation tend to be similar to the ones that are liable to offshoring (Blinder 2009, Blinder & Krueger 2013). A number of the theoretical and empirical studies engage with the issue of offshoring as an alternative hypothesis to technological change, most notably Goos et al. (2014) and Autor & Dorn (2013). Both find relatively minor effects of offshorability on occupational employment in Europe and the US, respectively. Moreover, the issue is complicated by the fact that offshoring is itself to a significant degree enabled by technological change (Goos et al. 2014). This is also highlighted by Grossman & Rossi-Hansberg (2008) who show that the productivity effect of offshoring is analogous to factor-augmenting technological change. While this does not mean that offshoring can simply be subsumed under technological change, it does explain why offshoring appears to lose out in a direct comparison with technological change, such as in Goos et al. (2014).

Another facet of globalisation is growing international trade, as shown in figure 1.4. Autor et al. (2013), Autor, Dorn & Hanson (2015) find significant effects of Chinese import competition on manufacturing employment in US commuting zones, but their analyses show that this effect is distinct from the technology effect. Whereas RBTC leads to a polarising reallocation of routine workers to other occupations, import competition reduces aggregate employment (Autor et al. 2015). Based on these studies, globalisation in its various manifestations does not appear to undermine the RBTC narrative as such. However, the findings serve as a reminder that RBTC is no catch-all explanation for employment and wage trends, and its effects must be carefully distinguished from those of other, parallel trends.

Moreover, like with SBTC, it is likely that RBTC theory has a “half-life” as new technologies enter the mass adoption phase and alter the way in which technology interacts with labour markets. Even if ICT is considered a 'general purpose technology' in the sense of Bresnahan & Trajtenberg (1995) and Aghion, Howitt & Violante (2002), it may be secondary inventions that have the greatest impact on the organisation of labour markets. It is naturally difficult to predict which new technologies (or new applications of existing technologies) will next reach the mass adoption stage (Mokyr et al. 2015), and what their implications for labour demand will be. Two studies by Beaudry, Green & Sand (2014, 2016) contain a first hint that the demand for skilled workers may (have) hit a plateau. Similarly, Acemoglu & Restrepo (2018) raise the prospect of high-skill automation which could significantly alter the dynamics that have underpinned the RBTC model. It is therefore important to reiterate that RBTC describes the relationship between technology and labour demand in capitalist economies of a certain level of development, and not any inherent qualities of technology.

2.3 The role of labour market institutions

By adhering to a supply-and-demand framework, the RBTC literature has largely neglected labour market institutions and the associated differences in the relative bargaining power of workers that research has shown influences some of the same outcomes. This section therefore surveys the research on labour market institutions in the context of declining routine occupations.

Labour market institutions are widely recognised for their crucial impact on various labour market outcomes (Checchi & García-Penalosa 2008, Nolan & Valenzuela 2019). LMI can be defined as "a system of laws, norms, or conventions resulting from a *collective* choice and providing constraints or incentives that alter *individual* choices over labor and pay" (Boeri & van Ours 2008, p. 3). As such, they have been the subject of many empirical studies in economics, sociology, and related disciplines, in particular with regard to their role in determining the level of income inequality in a country. Research on the relationship between LMI and occupational employment, technology adoption, and wage inequality is less prominent, however.

The effect of LMI is generally considered to operate by strengthening the relative bargaining power of workers. This reasoning is based on power resource theory, which is central to the institutional literature (see, e.g., Esping-Andersen 1990, Korpi 1983, 1985). Korpi (1985, p. 33) defines power resources as "the attributes (capacities or means) of actors (individuals or collectivities), which enable them to reward or to punish other actors." Crucially, "power resources can have important consequences even without being activated", so the mere theoretical possibility of, say, industrial action or recourse to an employment tribunal strengthens workers' position at the bargaining table (Korpi 1985, p. 33). Based on this logic, higher unionisation, stricter employment protection, a higher minimum wage, and generous replacement rates have all been linked empirically to more compressed earnings distributions (Ebbinghaus & Visser 1999, Esping-Andersen 1990, Gallie 2017, Koeniger et al. 2007, OECD 2015, Rueda & Pontusson 2000).

In this section, I analyse which insights the institutional literature can provide regarding employment change, technology adoption, and wage inequality. In this discussion, I mainly focus on unions and employment protection legislation, as these are the LMI that I engage with empirically in chapter 7, as well as the minimum wage, which is the best-researched of the major LMI.

2.3.1 Employment change and technology adoption

The relationship between LMI and which jobs exist in an economy is surprisingly understudied. Researchers have investigated the impact of unions on unemployment and insider – outsider divisions, but evidence on whether across countries and over time, more unionised occupations or sectors experience differential employment growth, is rare. Similarly, the relationship between unions and the adoption of new technologies is still unclear. EPL has been analysed with a view to aggregate employment effects, but evidence on occupation- or sector-specific effects is absent.

Unionisation

Based on the premise that unions confer a wage premium on their members,[22] power resource theory suggests a trade-off between unionisation and relative employment in routine occupations. One recent study that investigates this relationship directly is Parolin (2021). He hypothesises that insofar as unions are successful at protecting relative earnings, highly unionised occupations should face a faster contraction of their employment share. The empirical analysis of a large sample of American workers from the CPS between 1983 and 2017 finds that this was the case, as counterfactual estimations suggest that if union coverage had remained constant, the decline in routine intensity of employed workers would have been 21 percent steeper. Apart from this study, however, recent direct – or comparative – evidence regarding a trade-off between wages and employment for routine occupations is difficult to come by.

There are a number of older studies that find an overall negative employment effect of unions, without investigating routine occupations specifically (Leonard 1992, Nickell & Andrews 1983). This is linked to the argument that part of the labour market inequality-reducing effect of unions is due to marginal workers being pushed out of the workforce (Blau & Kahn 2002). However, recent studies also call into question the negative employment effect of unionisation. Brady et al. (2013) find no relationship between union membership and the probability of being employed in the United States between 1991 and 2010. Thus, there is some suggestive evidence that unionisation may reduce relative employment in routine occupations, assuming that routine workers are more likely to be unionised, but no strong conclusions are possible.

Another relevant question for the research questions in this thesis is whether and to which extent unions incentivise the adoption of labour-saving technology. Again, the literature does not provide a clear answer to this question. As Parolin (2021) discusses, it is not clear whether firms are more or less likely to invest in new

[22] This assertion is relatively uncontroversial, although recent research from the US also emphasises spillovers to non-union members (Brady, Baker & Finnigan 2013, Denice & Rosenfeld 2018, Wilmers 2017).

technologies when organised labour is stronger. Furthermore, even if technology adoption were faster in unionised firms, it is unclear if this would be beneficial for routine workers. The high labour costs due to unionisation may well be what drives technology adoption, with unions then fighting to implement the technology in a way that protects their employment and earnings (Fernandez 2001). On the other hand, Cardullo, Conti & Sulis (2015) find that across manufacturing industries in OECD countries between 1980 and 2000, union power reduced investment per worker, particularly in capital-intensive industries. Hence, the influence of organised labour on investment and technology adoption remains a relatively poorly understood phenomenon.

Employment protection legislation

This is true also of employment protection legislation. EPL refers to the rules governing the hiring and firing process, as well as the enforcement of these rules (Boeri 2011, OECD 2004). This is closely related to unions, since insurance against dismissal is a core function of unions. However, EPL extends to all workers covered by the legislation, regardless of their union membership status. Like unions, EPL improves the relative bargaining power of workers, by increasing the cost of hiring and firing workers. Effectively, EPL thus amounts to an "employer-borne tax on employment adjustment" (OECD 2004, p. 65).

A substantial literature discusses the relationship between EPL and labour market performance, but I am not aware of any empirical evidence regarding effects specifically on routine employment and technology adoption. In general terms, stricter EPL has been linked quite consistently to increased unemployment and reduced employment according to MacLeod (2011). For example, Autor, Donohue III & Schwab (2006), Autor, Kerr & Kugler (2007) and Bird & Knopf (2009) find that the adoption of wrongful-discharge laws in American states reduced employment, mainly for less educated workers. Feldmann (2008, 2009) finds a similar negative relationship in two large cross-country samples using EPL data from the Economic Freedom of the World index.

This is, however, contested by some more recent studies. In particular, studies of labour market liberalisation most often do not find the increase in employment that was used to justify the reforms (see, e.g., Arestis, Ferreiro & Gómez 2020, Kahn 2010). Instead, studies of liberalisation often find that the labour share declined due to the reforms (Ciminelli, Duval & Furceri 2018, Damiani, Pompei & Ricci 2020). This leads to the somewhat paradoxical situation that EPL is linked to lower employment and higher unemployment, but reductions in EPL do not appear to increase employment, at least not in permanent jobs. This may have to do with the time it takes for effects to materialise. A comprehensive review by Boeri, Cahuc & Zylberberg (2015) concludes that there are conflicting short- and long-term effects on employment. Moreover, Boeri et al. (2015) find that unskilled and inexperienced workers are more immediately affected by a relaxation of EPL, while the impact on skilled and experienced workers takes longer to materialise. Lacking, however, are studies of EPL and technology adoption, although to the extent to which EPL increases labour costs, the same arguments as for unionisation should apply.

Minimum wages

A substantial literature investigates the relationship between minimum wages and technology adoption, but since for data reasons my analyses in later chapters will not consider the minimum wage, I will only comment briefly here. Some research suggests that a high minimum wage might accelerate the substitution of machines for workers. Lordan & Neumark (2018, p. 51) find that "increasing the MW decreases significantly the share of automatable employment held by low-skilled workers." To illustrate, a minimum wage hike in Seattle, from USD 9.47 in 2014 to USD 13.00 in 2016, was associated with a 9 percent drop in hours worked in low-wage jobs and a concurrent 3 percent increase in hourly wages, amounting to a negative effect on total payroll (Jardim, Long, Plotnick, van Inwegen, Vigdor & Wething 2017).

The evidence concerning the aggregate employment effect of the minimum wage is far from clear, however. In the US, an influential study by Card & Krueger (1994) found no effect of a minimum wage hike in New Jersey on employment in the fast-food industry compared to neighbouring counties in Pennsylvania. This study ignited a long and inconclusive debate, with scholars such as Lordan & Neumark (2018) and Neumark, Schweitzer & Wascher (2004) arguing that there is a negative effect while Dube, Lester & Reich (2010), Dube & Lindner (2021) and Cengiz, Dube, Lindner & Zipperer (2019) find no evidence for employment losses. Equally in Germany and the UK, there is no evidence for substantial declines in employment after the introduction of the respective minimum wages (Ahlfeldt, Roth & Seidel 2018, Bossler & Gerner 2016, Dickens, Machin & Manning 1999, Dustmann, Lindner, Schönberg, Umkehrer & Berge 2021, Machin, Manning & Rahman 2003).

In conclusion, the evidence for effects of LMI on routine employment and technology adoption is surprisingly thin. Even for the minimum wage, which has received most attention, no strong conclusions appear possible. With regard to unions, a recent study by Parolin (2021) finds competing wage and employment effects of unionisation for routine workers, but evidence for an effect on technology adoption is thin. EPL has not been thoroughly investigated for possible effects on employment and technology adoption.

2.3.2 Wage inequality

Minimum wages

Research on LMI and wage inequality has been more extensive and generally concludes that LMI reduce inequality. Again, the minimum wage has received much of the attention. By acting as a wage floor, it is intuitively obvious that the minimum wage should reduce wage inequality in the lower parts of the wage distribution (Brown 1999, Dickens et al. 1999); minimum wage hikes can furthermore contribute to wage polarisation by raising low wages compared to medium wages.

In the US, the decline in the real value of the minimum wage has been offered as a leading alternative explanation for the increase in wage inequality after 1980 (Card & DiNardo 2002, DiNardo, Fortin & Lemieux 1996, Lee 1999). For example, Lee (1999) estimates that the falling federal minimum wage explains more than the entire rise in lower-tail inequality between 1979 and 1988, although Autor, Manning & Smith (2016) find a more modest yet still sizeable effect in the range of 30 to 40 percent of the observed increase. In the UK, Machin et al. (2003) and Dickens & Manning (2004) likewise find that the reintroduction of the minimum wage compressed the lower parts of the wage structure. The cross-country evidence is also clear that the minimum wage is associated with reduced lower-tail wage inequality (Koeniger et al. 2007, Pontusson et al. 2002). The main question is therefore not if the minimum wage reduces inequality, but by how much. With regard to inequality between routine workers and other groups, since the existing research looks at (regions of) the overall wage distribution, the impact would depend on which groups are covered by the minimum wage.

Unionisation

When it comes to wage inequality, the literature is less focused on the minimum wage, however. It is well established across social science subdisciplines that union-covered workers earn more, on average, than non-covered workers, in what is called the union wage premium (Freeman & Medoff 1984, Visser & Checchi 2011). Unionisation is therefore seen as a means to reduce overall inequality, besides other benefits such as increased employment security (Blanchflower & Bryson 2004, Rueda & Pontusson 2000, Salverda & Checchi 2015, Visser & Checchi 2011). By bargaining on behalf of their members and backed up by the threat of collective action, unions represent perhaps the most intuitive example of power resources in action (Korpi 1985).

This is illustrated in the aforementioned study by Parolin (2021) who investigates the impact of declining union membership on the employment and relative wages of routine workers in the US between 1983 and 2017. He finds that higher union coverage in a state-industry inhibits the decline in earnings returns to an occupation's routine-intensity by a sizeable margin. This, argues Parolin (2021, p. 941), "question[s] the extent to which automation deserves credit for the declining relative earnings of routine occupations" and underscores the importance of worker bargaining power.

That unionisation is generally and robustly associated with reduced inequality is established in numerous empirical studies. For example, VanHeuvelen (2018) shows that union decline has been central to understanding the increase in US wage inequality between 1973 and 2015, while Card et al. (2004) show evidence for Canada, the UK, and the US, and also point to differences between women and men. Koeniger et al. (2007) and Salverda & Checchi (2015) provide evidence from cross-country regressions that unionisation reduces wage inequality.

Some scholars highlight additional mechanisms by which unions may reduce wage inequality. For example, an influential study by Western & Rosenfeld (2011) argues that unions helped institutionalise norms of equity, to the benefit even of non-union workers. Thus, the authors show, the decline of private sector unionisation in the United States from 1973 to 2007 can account for up to one third of the growth in wage inequality. In contrast to power resource theory, Western & Rosenfeld (2011) attribute the impact of unions to the idea of a "moral economy" in which unions shape cultural, political, and institutional norms. This approach, which complements the rationalist theory of union threat, has also been adopted by VanHeuvelen (2018, 2020), while Wilmers (2017) finds that union activism, rather than market position alone, induces employers to raise wages.[23] The literature does not comment much on the stratifying effect of different patterns of union membership, although for example Card et al. (2004) discuss differences in unionisation rates between women and men and between different regions of the wage distribution.

Employment protection legislation

EPL has likewise been linked to lower overall wage inequality, although the volume of empirical research has been rather limited. Koeniger et al. (2007) study a panel of 11 countries for the period 1983 - 1998 and find that out of a range of institutional characteristics, a standard deviation reduction of EPL is associated with the largest

[23]By contrast, Breda (2015) argues that union wage premia are predominantly the result of rent extraction, with greater willingness of workers to organise and higher wage premia in high-rent firms. This does not seem wholly convincing as an explanation for the consistent and large effects observed in other studies, but regardless of the mechanism, the available evidence strongly suggests that unions have direct financial benefits for members while at the same time reducing the overall dispersion of wages.

increase in the 90/10 wage ratio, by 18 percent. As employment protection was indeed scaled back in many countries, this suggests that labour market liberalisation may have contributed significantly to the increase in inequality. Salverda & Checchi (2015) also find a negative association of EPL with various measures of the Gini coefficient on wages in a sample of 23 OECD countries. van der Wiel (2010) fleshes out the underlying mechanism by showing the causal effect of an exogenous extension of the term of notice on the wages of Dutch workers. She shows that this increase in firing costs led to an increase in hourly wages by 3 percent, after accounting for tenure and age. These findings are consistent with the power resource view of EPL. In particular, van der Wiel (2010) presents suggestive evidence that bargaining power, rather than investment in hard-to-fire workers, is driving the results.

Conversely, Bird & Knopf (2009, p. 211) note, "the reduced efficiency arising from discharge protections may cause demand of labor to drive inward, thereby forcing wages downward overall." Evidence for this, however, is very limited. Fishback & Kantor (1995) find that employers were able to pass on part of the cost of increased post-accident compensation in the form of lower wages. Autor et al. (2006) in their analysis of the effect of wrongful discharge laws in the US find robust employment effects, but no impact on wages. Due to these two opposing effects, the aggregate effect of EPL on wage levels is unclear, although the evidence is tilted in favour of a positive wage effect. The impact on inequality depends in either case on which workers are most affected by the policy. Such differential effects of EPL have not yet been investigated in detail. Van der Wiel's (2010) evidence suggests that lower educated workers benefit most from increased EPL, but her findings are far from conclusive.

Overall, the discussion of the evidence on labour market institutions has shown (1) the fundamental importance of LMI for interpreting labour market trends, (2) the crucial role for power resource theory in theorising the impact of institutions, and (3) the lack of a nuanced understanding of the effect of various LMI on outcomes in specific groups of occupations, such as high-routine occupations. Given that routine occupations are at the heart of the competing paradigm of RBTC, it would be crucial to develop specific hypotheses how LMI affect these occupations. This would allow researchers to better disentangle the effects of technology and institutions and elevate the debate about the relative importance of the different factors. The analyses in chapter 7 of this thesis contribute to this goal.

2.4 Key gaps in the theoretical literature

I have already at length analysed the limitations of RBTC theory with respect to its applicability across the developed world and the significant gaps in the LMI literature regarding outcomes such as relative employment, technology adoption, and relative wages. Rather than repeating these limitations, this section points out additional significant gaps in the theoretical literature concerning both RBTC and labour market institutions: the somewhat cavalier approach in much of the literature to measuring routine tasks, the simplistic equation of routine tasks with medium-wage tasks, and the lack of integration of the RBTC and institutional approaches.

In light of the preceding discussion, I briefly highlight why these limitations matter. First, instead of independently theorising what matters for capturing occupational tasks, scholars have, with few exceptions, been content to rely on "off-the-shelf measures" (Mahutga et al. 2018). Thus, nearly the entire RBTC literature is based on the measures of routine task intensity of Autor et al. (2003) and Autor & Dorn (2013). However, these measures are quite rudimentary, using a small set of questions that do not identify the key features of routine intensity, repetitiveness and codifiability, very well (Fernández-Macías & Hurley 2017). They are also based exclusively on outdated data from the US, despite variation in task content both between countries and over time. In chapters 3.1 and 5, I engage with this cavalier approach to measurement.

Secondly, comparative research often does not sufficiently account for country differences. This chapter has shown that scholars investigating RBTC outside the US tend to assume without verifying that routine occupations cluster near the median of the wage distribution there, too. This leads to a needlessly restrictive set of predictions which often do not correspond to empirical observations (Fernández-Macías & Hurley 2017, Goos et al. 2014). I argue in chapters 3.2 and 6 that there is no *a priori* reason to make this assumption. Taking country differences seriously can reconcile the observations that recent technological change is biased against routine occupations and that, while some countries have experienced polarisation, in many employment change has been upgrading. Thus, a second major shortcoming of the theoretical literature can be alleviated.

A third gap in the theoretical literature is a lack of efforts to integrate the RBTC and institutional approaches. Of course, this is not due to a lack of scholars' awareness that technology and institutions both matter (see, e.g., Antonczyk et al. 2018, Fernández-Macías 2012, Lemieux 2008, Nolan, Rahbari, Richiardi, Valenzuela Riviera & Nabarro 2017, Oesch 2013). For example, Fernández-Macías (2012) and Oesch (2013) argue that institutions cannot be ignored when analysing technological and labour market change, and Nolan et al. (2017) state that neither an "Economics First" nor a "Politics First" approach does the complexity of labour market trends justice. Yet, most research has proceeded within disciplinary boundaries with little regard to these findings. This may have to do with the fact that both literatures focus on slightly different outcome variables, as chapter 2.3 has shown. I address the lack of integration in the existing theoretical literature in chapters 3.3 and 7.

2.5 Conclusion

This chapter has provided a comprehensive overview of the theoretical underpinnings of this thesis. First, I have in detail explored the evolution of the RBTC model out of the SBTC model, its key tenets and empirical implications. The model is shown to convincingly explain employment and inequality trends in the United States over the past three decades. I then reviewed the evidence for RBTC across the developed world, which turns out to be significantly less conclusive than for the US. Finally, I briefly assessed the evidence for globalisation as an alternative explanatory paradigm, but found that it cannot seriously challenge the RBTC narrative. Even as I go on to argue that some assumptions of the RBTC model must be modified in order for the model to yield generalisable predictions, the RBTC framework remains the backbone of this thesis. However, I argue that sociologists and other scholars have made important contributions that scholars of RBTC cannot ignore.

The second theoretical paradigm that I engage with in this thesis is the institutionalist perspective which argues that labour market trends, especially pertaining to inequality, are best explained by political choices over labour market institutions. Since my interest is in the interrelationship of RBTC and institutions, the discussion focused specifically on the impact of unionisation, employment protection, and minimum wages on the relative employment and wages of routine occupations as well as on incentives for technology adoption. Here, the most striking finding is a severe lack of theoretical and empirical research on the impact of these LMI on specific occupational groups, even as effects on overall inequality, for example, are well documented.

Following this analysis, a final section pointed out some further gaps in the theoretical literature such as the insufficient theorisation of routine tasks, the simplistic equation of routine tasks with medium-wage tasks, and the lack of integration of the RBTC and institutional approaches. These limitations are addressed in the remaining chapters of this thesis.

3
Contributions to the literature and analytical strategy

Contents

This chapter outlines the main theoretical contributions of this thesis and the overall analytical strategy. Based on the discussion in chapter 2, I lay out the contributions of this thesis to a more comprehensive understanding of the effects of biased technological change and labour market institutions on labour market outcomes based on the three main research questions: How can the measurement of technological change be improved, how can we explain country differences in outcomes, and do robotisation and labour market institutions drive these differences? In doing so, I also sketch how the different empirical chapters relate to one another to form a coherent whole. Following this overview of the content of the thesis, I provide some detail on the analytical strategy used to arrive at its conclusions.

3.1 Sound foundations: Rethinking measurement

I argued in chapter 2.4 that the indices used in much of the literature on biased technological change lack a thorough justification of the underlying assumptions, data, and measures. In light of this finding, it is all the more important to measure occupational task content in a way that corresponds to the underlying theoretical concepts and that accounts for the relevant dimensions of variation, as research question 1 demands: robust social scientific insights require solid underlying data. In chapter 5, I present a set of measures that satisfy these criteria, using data from the EWCS for 27 EU countries between 2000 and 2015.

Like other scholars before me, I conceptualise routine tasks as activities that are repetitive and/or codifiable (Acemoglu & Autor 2011, Autor et al. 2003, Fernández-Macías & Hurley 2017). My approach is distinct, however, for three main reasons. The first, and most important, is the detailed attention to the congruence of concept and measurement. Even though Fernández-Macías & Hurley (2017) have previously argued that existing measures do not correctly identify routine tasks based on the measures used, my contribution is the first to perform a thorough analysis of the construct and criterion validity of the proposed measures and thus present quantitative evidence to back up my theoretically informed choices of variables. Secondly, the measure of routine intensity is accompanied by the first index of task complexity that approximates susceptibility to SBTC in a task-based framework, enabling the analysis of SBTC alongside RBTC.

The third key feature that distinguishes my measures from other approaches, is that they provide new opportunities for research into task change over time and differences between countries. These crucial dimensions of variation have been largely ignored in previous research, even though theoretical considerations suggest that differences may not be trivial. I thus provide a body of descriptive evidence that shows systematic changes over time as well as sizeable country differences in routine task intensity and task complexity.

Overall, chapter 5 makes an important theoretical and empirical contribution. Based on a comprehensive theoretical discussion, it provides a set of measures that serve as the empirical backbone for my analyses in subsequent chapters that differ in important aspects from the prominent measures in the literature. It also yields important empirical findings, by showing how occupational tasks have changed over time and differ between countries. Thus, by developing sound foundations for analyses in this thesis and beyond, chapter 5 makes a vital contribution to the theoretical understanding of the role of occupational tasks and the empirical knowledge about their distribution.

3.2 One size fits not all: Incorporating diversity

The second research question asks whether the diverse patterns of employment change in developed countries can be reconciled with the key tenet of RBTC, that technological change everywhere substitutes for routine workers? Due to the heavy reliance on the American experience in the formulation of the theory, many economists have assumed, implicitly or explicitly, that routine occupations are always medium-wage and thus RBTC would play out similarly in other developed countries, generally resulting in polarisation. However, in chapter 2 I showed that polarisation has by no means been a universal finding in studies of employment change (Fernández-Macías 2012, Oesch 2013). From this it follows that either technological change is not always routine-biased, or routine occupations are not always medium-wage. Of course, the former would contradict the central tenet of RBTC theory. This paradox constitutes a major gap in the literature which I address in chapter 6.

My novel theoretical contribution in this chapter is the proposal of a refined model of RBTC that reconciles the key insight that the routine intensity of an occupation determines whether it is complemented or substituted by technology with the observation that occupational upgrading has continued in many countries (Eurofound 2017, Oesch 2013).

The central element of this refined model is that polarisation depends on high-routine occupations being concentrated around the median of the wage distribution. This is documented in numerous articles on the US (Autor & Dorn 2013, Autor et al. 2008), but not in many of the countries which the RBTC framework has been applied to, for example by Goos et al. (2014). Treating the position of routine occupations in the wage hierarchy as an empirical question implies that RBTC need not lead to polarisation. If the routine occupations that are being replaced by technology are low-wage, RBTC may contribute to upgrading instead.

I develop a simple visual technique called task-wage curves with which the relationship between average occupational wages and characteristics such as routine intensity can be assessed. I fit a locally weighted fractional polynomial function on the unconditional relationship between the variables of interest. This allows for straightforward country comparisons of the relative position of routine occupations in the wage distribution. Analysing the correspondence of patterns of employment change with different shapes of the routine-wage curve then allows me to test the predictions of my refined theory.

Indeed, the finding is striking: where routine occupations do cluster around the median wage, occupational polarisation is common, but where they are closer towards the bottom of the wage hierarchy, upgrading is the norm. This indicates that the restrictive assumption that routine occupations are generally medium-wage has prevented the RBTC model from being convincingly applied in many cases. Thus, by taking seriously diversity between countries, the refined model of RBTC proposed in chapter 6 makes it possible to significantly broaden the scope of RBTC theory.

3.3 Digging deeper: Wage hierarchies explained

The third research question builds on the finding from chapter 6 that patterns of relative occupational wages correspond to patterns of employment change. It therefore asks whether these wage patterns are the result of heterogeneous effects of robotisation, LMI, and possible interactions between the two on different occupational groups. By pursuing this question, chapter 7 attempts to fill a gap in our knowledge about the determinants of occupational wage hierarchies, and shed light on the correlates of wage inequality within the lower two thirds of the wage distribution. In the process, it also qualifies the sociological finding that occupational hierarchies are largely invariant across countries (Hout & DiPrete 2006, Treiman 1977).

This chapter is like previous chapters informed by RBTC theory, but the analytical focus shifts towards the power resource framework that is used in the sociological and social policy literatures (Korpi 1983). I propose a theory in which occupational wage differentials are determined by the relative productivity and relative bargaining power of occupations. This approach entails two theoretical innovations. Firstly, while the power resource framework has commonly been used to analyse the relative power of workers across the wage distribution (using measures such as the 90/10 wage ratio), in this model two distinct groups of occupations are compared based on theoretical criteria. Secondly, by investigating interactions between robots and LMI, I am the first to really integrate the productivity and power resource mechanisms.

The chapter also contributes a methodological innovation in the form of the manufacturing wage premium, which allows me to express the key feature of the RTI-wage curves numerically. Routine manufacturing occupations are considered the most exposed to RBTC and at the same time ordinarily enjoy a wage premium compared to other non-professional occupations. Thus by construction, a high manufacturing wage premium corresponds to a hump-shaped routine-wage curve in chapter 6. This measure can therefore serve as the dependent variable in the analysis of research question 3. Comparing two theoretically defined groups of occupations is arguably preferable to using an overall measure of inequality such as the 50/10 wage ratio, since the outcome measures can be tailored to the research question at hand and can thus be interpreted more straightforwardly. In the context of the RBTC literature, this approach constitutes another novelty.

Based on these theoretical and methodological innovations, chapter 7 contributes the first comparative investigation of wage inequality between routine manufacturing and other non-professional occupations. It also offers the first analysis of the role of robotisation alongside and in dependence on LMI and thus extends the research of Graetz & Michaels (2018). While the directional hypotheses which I develop based on the robotisation and welfare state literatures are mostly not supported, the study still confirms that robotisation and LMI do not affect all workers equally, and that moreover the effect of robotisation on wages may depend on institutional characteristics (Eichhorst & Marx 2015).

Thus, by showing that small differences in occupational wages hierarchies matter, Hout & DiPrete (2006) and Treiman (1977) notwithstanding, and that robotisation and LMI may have heterogeneous effects on different occupational groups, in line with an extension of the power resource literature, chapter 7 makes an important contribution to a better understanding of the interplay between technology and institutions with regard to labour market outcomes. The three empirical chapters together form a coherent whole and contribute significantly to answering the overall research puzzle, how technological change and LMI affect employment, wages, and inequality.

3.4 Analytical strategy

After clarifying the substantive contributions of my thesis, the remainder of this chapter discusses the overall analytical strategy and common methodological themes that pervade the empirical chapters. Naturally, the detailed research questions each require their own tailored analytical approach. Nevertheless, there are some common methodological choices that merit a brief discussion to spell out not only what my thesis contributes, but also how it does so. Hence, in this section, I briefly comment on the statistical approach, level of analysis, and approach to sample selection throughout the thesis.

3.4.1 Statistical approach

For the empirical analyses in this thesis, a variety of statistical techniques are employed. As my analyses often challenge widely held assumptions, descriptive techniques play an important role in establishing the research puzzle, closely analysing the available data, and pointing towards possible answers to the research questions (Winkler 2009). Alongside tailored and innovative descriptive analyses, regression models are employed to investigate multivariate relationships and the generalisability of my arguments.

As I explained above, I introduce two new concepts into the literature: task-wage curves and the manufacturing wage premium. The task-wage curves visualise the relationship between average occupational wages and task characteristics such as routine intensity. The manufacturing wage premium, on the other hand, compares the average wages in high-routine occupations that are predominantly in the manufacturing sector with the average wages in medium- to high-routine occupations that are not predominantly in the manufacturing sector. These new concepts are at the heart of my descriptive analyses in chapters 6 and 7 and allow me to present my insights in a clear and straightforward manner. In chapter 5, I mainly rely on the descriptive analysis of dispersion and correlations to argue in favour of my alternative indices of task content.

For the regression analyses, I rely on OLS panel models. This simple yet effective approach is the most suitable method for the often comparatively small (in chapters 6 and 7) or heavily unbalanced (in chapter 7) panels in my analyses. In doing so, I follow several other scholars in the field who use a similar approach, such as Goos et al. (2014). I use fixed effects to capture country- and occupation-specific variation and clustered standard errors (Abadie, Athey, Imbens & Wooldridge 2017). Deviations from this general approach are discussed in the respective chapters. Given this statistical approach, the findings of my multivariate models should not be interpreted as causal effects. Rather, they point out associations with employment, wages, and inequality that are indicative of a potential causal relationship. Statements regarding the "effect" of the independent variables should be interpreted in this vein.

3.4.2 The level of analysis: 2-digit ISCO-88 occupations

In chapters 5 and 6, my analyses are at the level of 2-digit ISCO-88 occupations.[1] The ISCO occupational classification has four levels: 9 major groups are subdivided into 27 sub-major groups, which in turn comprise 115 minor groups and 389

[1]For chapter 7, I compute country-level measures based on 2-digit occupational and wage data.

unit groups.[2] The sub-major groups are listed in the appendix. Compared to some other articles which use 3-digit occupational codes, this obviously entails a loss of precision (see, e.g., Autor et al. 2008, Cortes 2016, Oesch & Piccitto 2019). However, the sample sizes of some LIS country-years make it infeasible to conduct the analysis at that level.

Moreover, the literature suggests that the 2-digit level is sufficiently detailed for the comparative analyses in this thesis. One important reason for this is that the task measures are likewise at the 2-digit level. This means that even if the employment data were 3-digit, the gain in precision in estimating the relationship between occupational tasks and employment or wages would be very limited. With this, I follow the seminal article of Goos et al. (2014) who also employ a routine intensity measure that is aggregated at the 2-digit level. Moreover, in their 2007 article on polarisation in the UK, although they show results for 3-digit SOC90 occupations, Goos & Manning (2007) report that their results are robust to the level of disaggregation chosen. This suggests that substantial differences between 2-digit and 3-digit analyses are unlikely.

The same applies to whether or not occupations are cross-classified by industry: the occupation component appears to be the much more important one (Goos & Manning 2007). Similarly, Acemoglu & Autor (2011), in what is still one of the definitive accounts of the labour market implications of technological change, report that in the US, within-industry shifts in occupational employment by far outstrip the size of between-industry effects. Fonseca et al. (2018) find the same for Portugal. In Goos et al. (2014), the within- and between-industry shifts are more evenly sized, but in any case the effects point in the same direction. Therefore, there appears to be no danger of misrepresenting employment trends by focusing on occupational variation. For these reasons, it seems appropriate to conduct the analyses with both the task and employment/wage data at the level of 2-digit occupations.

[2] In practice, I operate with 26 sub-major groups: group 62 (subsistence agricultural and fishery workers) is subsumed under 61 (market-oriented skilled agricultural and fishery workers), since code 62 does not exist in most countries, and where it does, the number of individuals employed in the occupation is negligible.

3.4.3 Sample selection

In chapter 5, I develop measures of occupational task content based on EWCS data from the EU-27 and EU-15 country groups. Using data from a well-defined geographical region and a single data source for the creation of the task indices ensures the consistency of the indices. However, the samples of the EWCS, the LIS, and the International Federation of Robotics (IFR) robot database do not always coincide. Hence, the nature of my data sources necessitates the use of different samples, and the analyses in chapters 6 and 7 each use different samples from chapter 5. In this section, I explain in more detail my sample selection strategy and discuss potential concerns.

The analyses in chapter 6 pertain to a 10-country sample that includes Chile, Czech Republic, Finland, Germany, Hungary, Luxembourg, Netherlands, Slovenia, Spain, and the United States. The sample selection strategy is based on data availability in the LIS and includes two criteria:

1. The country is included in each LIS wave from wave IV (1995) to IX (2013).

2. The LIS reports consistent occupational and earnings data throughout this time.

Of the resulting sample, all countries except Chile and the United States are included in the EU-27 version of the task indices alongside 19 other countries; Finland, Germany, Luxembourg, The Netherlands, and Spain are included in the EU-15 task indices together with 10 additional (predominantly Western European) countries. Of course, perfect overlap between the two samples would be preferable, but because of the limited number of countries that consistently report occupational and earnings data to the LIS and the European nature of the EWCS, this is not possible. The 10-country sample thus represents a compromise between geographic and temporal coverage in the LIS. Moreover, as I explain in chapter 5, the task measures prove to be applicable out of sample as well, justifying in my eyes their application in Chile and the US.

The reasons for changing the sample composition again in chapter 7 are less obvious, but in my eyes no less compelling. For one, Luxembourg does not report robot data to the IFR and institutional information is not available for Chile in the Comparative Political Data Set (CPDS). Since the regression analyses in chapter 7 are at the country level rather than the country-occupation level, this reduced sample would be too small for reliable statistical inference. At the same time, an uninterrupted time series is less essential in these analyses, since the data are modelled as a panel with time fixed-effects owing to the small number of observations for most counties. Consequentially, expanding the sample to include all country-years for which the LIS, IFR, and CPDS provide all the relevant information, is the prudent choice. More details on this sample can be found in chapter 7 and the corresponding appendix.

Even though all countries in in my analyses are industrialised economies and members of the OECD, it is also clear that there remains a large degree of variation in terms of economic development and other relevant characteristics such as welfare regimes or political heritage. Table 3.1 provides some contextual information about the countries in the 10-country sample. All regions of Europe as well as Southern and Northern America are represented, and both population and GDP per capita vary widely. Moreover, half the sample consists of countries that had become democratic at most 20 years before the start of the period of analysis. In terms of welfare regimes, only Finland, Germany, The Netherlands, and the United States are included in Esping-Andersen's (1990) typology. Yet, all three regimes are represented. Less well-known typologies such as that of Andersen (2012) also suggest a relatively even distribution across regimes. This multifaceted diversity suggests that robust results in this sample should stand a good chance of holding up in a larger sample of developed countries.

Table 3.1: Basic information on the 10-country sample.

Country	Region	Population	GDP p. c.	Democracy	Welfare regime
Chile	SA	18.3	24,226	1990	
Czech Republic	EE	10.7	40,862	1989	
Finland	NE	5.6	48,668	1917	C (E-A); N (A)
Germany	WE	79.9	53,919	1945	C (E-A); C (A)
Hungary	EE	9.7	32,945	1990	
Luxembourg	WE	0.64	114,482	1919	C (A)
Netherlands	WE	17.3	56,935	1917	N (E-A); C (A)
Slovenia	SE	2.1	39,088	1991	S (A)
Spain	SE	47.3	40,903	1975	S (A)
United States	NA	335.0	62,530	1776	L (E-A)

Note: All data, except on welfare regimes, are from the CIA World Factbook (CIA 2021). Population in millions, GDP per capita in 2010 PPP USD. Geographical regions according to UN codes: SA – Southern America, EE – Eastern Europe, NE – Northern Europe, WE – Western Europe, SE – Southern Europe, NA – Northern America. Welfare regimes according to Esping-Andersen (1990) (E-A) and Andersen (2012) (A): N: social-democratic/Nordic; C: conservative/continental; L: liberal/Anglo-Saxon; S: Southern.

3.5 Conclusion

This chapter has provided an outlook on the contributions of this thesis to the social-scientific literature and on the analytical strategy I employ. Chapter 5 develops a refined set of measures of occupational task content which address the shortcomings in the existing literature and provide sound foundations for the further analyses. Furthermore, the chapter shows that occupational routine intensity and complexity have changed over time and differ between countries. Chapter 6 proposes a refined model of RBTC in which the relative wages of routine occupations may vary between countries. Thus, by incorporating diversity between countries, the model can accommodate different patterns of employment change and thus resolve the paradox of non-pervasive polarisation in the RBTC literature. Chapter 7 digs deeper and shows that robotisation and LMI cannot explain the wage patterns which have proven conducive to polarisation. There is, however, evidence for heterogeneous effects of robotisation and LMI on different occupational groups, and a moderating effect of EPL on the relationship between robotisation and relative wages.

I arrive at these conclusions based on the extensive use of descriptive statistical methods and the appropriate panel regression techniques. The analyses are performed predominantly at the 2-digit ISCO-88 level in a diverse sample of countries in Europe and the Americas. My research thus makes a significant contribution to the theoretical debates and to the corpus of empirical evidence concerning the effect of technological change and labour market institutions on employment, wages, and inequality.

4
Data

Contents

After the previous chapter detailed the theoretical and empirical contributions of this thesis and the analytical strategy, this chapter introduces the main data sources that are used throughout its empirical part.

4.1 European Working Conditions Survey: Task content data

The task content measures in this study are calculated from the waves 3 – 6 of the European Working Conditions Survey. The EWCS is a quinquennial survey conducted by the European Foundation for the Improvement of Living and Working Conditions (Eurofound). The purpose of the survey is, among other things, to "assess and quantify working conditions of both employees and the self-employed across Europe on a harmonised basis" (Eurofound 2021). To this end, in every wave a random sample of workers (employed and self-employed) is interviewed face to face.[1]

Waves 3 – 6 provide workplace information for the years 2000, 2005, 2010 and 2015. I work with samples comprising the EU-27 and EU-15 countries. The total N amounts to 107,488 individuals in the EU-27 sample and 74,895 in the EU-15 sample. Since the Eastern European enlargement of the EU only took effect in 2004, the affected countries are not included in the 2000 wave of the EWCS. Hence, to keep the samples as consistent as possible, I mainly work with the EU-15 data throughout the thesis. Table 4.1 below shows the sample sizes in each country that are used to calculate the indices. The precise construction of the routine and complexity indices is described in chapter 5.

[1]The target sample size is generally between 1,000 and 4,000 individuals depending on country size. There are a few exceptions: the target sample size was 500 in Luxembourg in 2000 and 600 in Cyprus, Estonia, Luxembourg, Malta and Slovenia in 2005.

Table 4.1: EWCS countries and sample sizes

Country	Analytical sample size				
	2000	**2005**	**2010**	**2015**	**Total**
Austria	1,288	791	868	946	3,893
Belgium	1,421	938	3,749	2,480	8,588
Denmark	1,361	916	1,025	966	4,268
Finland	1,318	977	955	954	4,204
France	1,440	972	2,869	1,489	6,770
Germany	1,343	895	2,023	1,948	6,209
Ireland	1,413	919	906	1,000	4,241
Italy	1,329	876	1,334	1,247	4,786
Luxembourg	432	531	927	956	2,846
Netherlands	1,415	996	984	1,005	4,400
Portugal	1,387	947	967	997	4,298
Spain	1,357	921	939	3,163	6,380
Sweden	1,417	984	913	958	4,272
United Kingdom	1,438	997	1,489	1,558	5,492
Total EU-15	**19,751**	**13,578**	**20,940**	**20,623**	**74,895**
Bulgaria	0	923	893	1,000	2,816
Cyprus	0	554	941	959	2,454
Czech Republic	0	752	886	927	2,565
Estonia	0	517	848	929	2,294
Hungary	0	963	967	957	2,887
Latvia	0	896	927	845	2,668
Lithuania	0	694	819	932	2,445
Malta	0	536	918	953	2,407
Poland	0	857	1,310	1,072	3,239
Romania	0	882	928	944	2,754
Slovakia	0	934	889	915	2,738
Slovenia	0	536	1,316	1,474	3,326
Total EU-27	**19,751**	**22,622**	**32,582**	**32,530**	**107,488**

Note: This table lists the countries included in the analyses based on the EWCS and the sample size in each EWCS wave.

4.2 Luxembourg Income Study: Employment and wage data

The main data source used in this thesis is the Luxembourg Income Study (LIS). It supplies the data on occupational employment and wages that are crucial for the analyses in chapters 6 and 7. Like the EWCS, the LIS sample includes information from employed as well as self-employed workers.[2] The LIS is the largest available income and occupational database of harmonised microdata collected from about 50 countries (Luxembourg Income Study 2021). The LIS operates with survey waves which are scheduled approximately every 3 years. However, annual data are available for an increasing number of countries. My analyses in chapter 6 are based on waves IV to IX, which cover the 1995 – 2013 time period, and focus on 10 countries: Chile, Czech Republic, Finland, Germany, Hungary, Luxembourg, Netherlands, Slovenia, Spain, and the United States. This selection is based on the availability in these countries of an uninterrupted time-series throughout the 6 LIS waves.[3] The sample in chapter 7 is similar but there are some differences due to data constraints. For example, Luxembourg does not report robot data and Chile lacks the institutional data and can therefore not be included in the main analyses. In turn, several other countries are included even if they do not provide an uninterrupted time-series in the LIS. This is necessary to achieve sufficient statistical power for the regression analyses.

Compared to other data sources, the LIS boasts several strengths that make it appealing for comparative researchers, foremost its geographical and temporal coverage. While other high-quality data sources such as the EU-LFS, EU-SES, or EU-SILC include only European countries, the LIS data allow me to also analyse non-European countries such as Chile and the United States. Especially the inclusion

[2]The inclusion of self-employed workers could lead to concerns about comparability due to differences between countries in self-employment rates. However, the seminal studies in the field, including Autor et al. (2003), Fernández-Macías (2012), and Goos et al. (2014), all include the self-employed. Thus, excluding the self-employed would reduce the comparability with other studies. This would make it harder to determine whether the proposed mechanisms are at play.

[3]Data for additional countries are now available for the 1995 – 2013 time period; however, for consistency with Haslberger (2021*b*), in this thesis, I stick to the original 10-country sample in chapter 6.

Table 4.2: Comparison of LIS and alternative data sources

	LIS	**EU-LFS**	**EU-SES**	**EU-SILC**
Time frame	~1980 - now	1983 - now	2002 - now	2004 - now
Countries	OECD +	EU +	EU -	EU +
Occupational data	✓	✓	✓	
Earnings data	✓		✓	✓

Note: EU-LFS: EU Labour Force Survey; EU-SES: EU Structure of Earnings Survey; EU-SILC: EU Statistics on Income and Living Conditions.

of the United States, on which so much of the existing literature on technological change focuses, is a major advantage. But also Chile, as one of the first non-Western OECD countries, promises to be an enriching addition to comparative datasets. Furthermore, other large-scale European datasets either do not include feasible earnings data (EU-LFS) or occupational data (EU-SILC), or cover too short a time period (EU-SES and EU-SILC). Thus, to use these datasets, it would be necessary to combine wage and occupational data from different sources, and even then, the available time period would be significantly shorter. With the LIS, on the other hand, employment and wage data pertain to the same individuals and span a period of almost 20 years. The comparison can be seen in table 4.2.

However, the LIS has some important disadvantages as well. For one, in a number of countries, individuals' occupations were not originally recorded using the ISCO-88 classification, or the new ISCO-08 classification was adopted around 2011. This makes it necessary to recode the occupational data to ISCO-88 in a number of countries. Although I follow the admirable and very careful effort by Mahutga et al. (2018), this will inevitably have introduced some measurement error.[4] Secondly, the LIS sample sizes are highly variable and, on average, small compared to sources like the EU-LFS. This can be seen in table 4.3. In Spain in 1995 and in Hungary throughout, the sample size is between 1,000 and 2,000, while it approaches 100,000 in Chile and the US in some years, with most country-years

[4] Based on the work of Mahutga et al. (2018), I have extended the number of recoded datasets as more LIS waves have become available and corrected some minor coding mistakes in the process.

Table 4.3: LIS samples for chapter 6

Country	Survey wave					
	∼1995	∼2000	∼2004	∼2007	∼2010	∼2013
CL	46,001	80,079	87,443	97,126	76,054	85,897
CZ	30,652	8,476	4,373	11,454	8,616	7,564
DE	6,484	11,886	10,907	10,345	16,012	13,708
ES	1,144	5,057	13,511	14,099	11,468	10,482
FI	9,144	11,997	11,966	12,195	9,829	11,610
HU	1,579	1,819	1,745	1,624	1,550	1,599
LU	2,624	2,737	3,934	4,362	6,020	4,265
NL	4,847	4,742	10,155	11,639	11,496	11,052
SI	3,064	5,208	4,682	4,687	4,685	4,071
US	67,989	101,879	96,077	95,475	89,275	61,859

Note: Samples used to compute the employment shares and occupational wages in chapter 6. Where annual data are available, only the year directly corresponding to the LIS wave is shown.

sitting around 10,000 observations. As I use only 26 ISCO 2-digit categories, even in countries with a lower sample size, the findings are generally reliable. Table 4.3 illustrates, however, why analyses at the 3-digit level would not be feasible with the available LIS data in this sample.

Perhaps the most serious drawback of the LIS is that it provides hourly wage data only in a handful of countries. This makes using the hourly wage variable infeasible in multi-country studies. Therefore, I follow Mahutga et al. (2018) and work with the variable for annual personal labour income (*pilabour*). This variable includes "cash payments and value of goods and services received from dependent employment, profits/losses and value of goods from self-employment, as well as the value of own consumption", and excludes capital and transfer income (LIS 2019). Thus, without information on hours worked, my estimates of average occupational wages may be conflated by differences between occupations and countries in average annual hours worked.

Hourly wage data would therefore be preferable; however, such detailed data are not available for many of the LIS country-years for which we have occupational data. Thus, to preserve the main advantage of the LIS which is its unmatched geographical and temporal coverage, I work with the annual income variable. To alleviate concerns that my findings could be driven by differences in hours worked, where appropriate, I conduct supplementary analyses with hourly wage data for countries for which such data are available, or investigate patterns of (female) part-time work. The results of these analyses generally indicate that my findings for annual wages are robust to considering hourly wages instead.

4.3 International Federation of Robotics: Data on industrial robots

Chapter 7 investigates the impact of robotisation on the relative wages of routine manufacturing workers. Data on industrial robots are taken from the International Federation of Robotics database, covering the period from 1993 - 2018. This dataset is the gold standard when it comes to the analysis of the aggregate impact of robots on jobs and is used for example by Dauth, Findeisen, Suedekum & Woessner (2019), Graetz & Michaels (2018), and Klenert, Fernández-Macías & Antón (2020). The details of my approach to measuring robot density are described in chapter 7.

4.4 Macroeconomic and institutional variables

My analyses furthermore account for a number of macroeconomic and institutional factors. The data are mainly taken from the Comparative Political Data Set (Armingeon et al. 2020). This dataset compiles a range of important country-level indicators such as union density, employment protection, and absolute employment by economic sector. It is also the source of my control variables relating to globalisation. Some missing observations are added to my dataset from the OECD and Statista databases. The detailed choices are discussed in chapter 7.

5

Rethinking the measurement of occupational task content

Contents

5.1 Introduction

Ever since the seminal article by Autor et al. (2003), routine-biased technological change has been the dominant explanation for employment and wage trends in the technological change literature. The prediction of the theory is that technology will substitute for workers in producing routine tasks, reducing the relative employment and in some scenarios the relative wage of workers in routine-intensive occupations.

Correctly identifying to which degree different occupations entail routine tasks is therefore essential for empirical analyses of this theory: a bias in the foundational measure would undermine all subsequent conclusions. To empirically investigate the predictions of the theory, Autor et al. (2003) developed a set of measures of occupational task content which have since, by virtue of academic primogeniture, enjoyed a near-monopoly in the task literature: important studies such as Autor et al. (2008), Goos & Manning (2007), and Goos et al. (2014) use this set of measures or a simplified version developed by Autor & Dorn (2013).

Only recently have scholars begun to subject these measures to greater scrutiny. Invariably, they find conceptual and empirical problems with the standard indices (Fernández-Macías & Hurley 2017, Handel 2017, Sebastian & Biagi 2018). This chapter adds to this recent literature. It poses the following research question: how can occupational task content be measured in a way that corresponds to the underlying theoretical concepts and that accounts for the relevant dimensions of variation? It argues that the congruence of theoretical concepts and their empirical implementation needs to be improved and that suitable indices need to account for differences over time and between countries. To this end, the chapter first subjects the task measures of Autor et al. (2003) as well as a recent alternative approach developed by Fernández-Macías & Hurley (2017) to a critical review with regard to their theoretical and empirical properties.

It then develops improved measures of routine task intensity and task complexity using workers' self-assessments in the European Working Conditions Survey. One illustrative improvement is the case of office clerks who based on the methodology of Autor et al. (2003) are characterised as the most routine-intensive occupation. This assessment completely disregards the evolution of this occupational group since the 1970s as it has computerised and clerks have taken over formerly managerial tasks. Using recent data and a more suitable set of variables, the approach proposed here places office clerks near the middle of the routine distribution, which is more in line with today's workplace requirements.

The second contribution of this chapter is to show that occupational tasks in the EU-15 countries have changed over time and vary between countries. This further strengthens the case that the measures of Autor et al. (2003) are not well suited for analysing recent employment changes in diverse samples. This chapter therefore answers the first research question of this thesis and makes an important methodological contribution which enables a more theoretically informed understanding of technological change. As such, it also forms the basis for the empirical analyses throughout this thesis, which rely on the insights and measures derived here.

5.2 Overview of existing operationalisations

5.2.1 What are routine tasks?

The RBTC approach explicitly asks "what it is that people do with computers" (Autor et al. 2003, p. 1280), and thus takes a first step away from the black-box view of technological change that has been criticised by sociologists (Fernandez 2001). In another seminal article in the task literature, a task is defined as a "unit of work activity that produces output" (Acemoglu & Autor 2011, p. 1118). Essentially, the RBTC argument predicts a reallocation of employment based on the task composition of occupations. The core of Autor et al's. (2003, p. 1280) argument is worth quoting in full and stipulates

> "(1) that computer capital substitutes for workers in carrying out a limited and well-defined set of cognitive and manual activities, those that can be accomplished by following explicit rules (what we term 'routine tasks'); and (2) that computer capital complements workers in carrying out problem-solving and complex communication activities ('nonroutine' tasks)."

Thus, production tasks are allocated to workers or capital based on comparative advantage in performing the respective tasks, where capital has an advantage in performing routine tasks and workers have an advantage when it comes to non-routine tasks. Elsewhere in the article, the authors argue that routine tasks "require [the] methodical repetition of an unwavering procedure" (Autor et al. 2003, p. 1283), thus introducing the element of repetitiveness. Other definitions require that routine tasks be "expressible in rules such that they are easily programmable and can be performed by computers at economically feasible costs" (Spitz-Oener 2006, p. 239), or define routine-intensity as "the extent to which an occupation is automatable or codifiable" (Caines, Hoffmann & Kambourov 2017, p. 302).[1] Hence, conceptually, the labour economics literature focuses on codifiability and repetitiveness as the distinguishing features of routine tasks.

5.2.2 Existing operationalisations of occupational task content

This chapter engages with two of the most influential approaches to quantifying occupational task content.[2] The first, adopted by most labour economists, follows the pioneers of the task-based approach. Autor et al. (2003) identified five task dimensions in their empirical analysis which Autor & Dorn (2013) have consolidated and formalised into a framework which classifies tasks as routine, abstract, or manual, and has become mainstream in the economics literature and beyond (see Sebastian & Biagi 2018 for an overview).

However, this dominance is not the result of rigorous debate on how to best measure occupational task content, but mainly of convenience turned convention. For example, Mahutga et al. (2018, p. 83) state as their rationale for using the Autor & Dorn (2013) measures that they “follow recommendations elsewhere that researchers use existing measures of the content of occupations as much as possible.”

The second approach, represented by recent studies in economic sociology, retains the basic logic of the task framework. However, its proponents choose different task dimensions, variables, data sources and units of analysis to operationalise task content (Fernández-Macías & Hurley 2017). Further key differences from the Autor & Dorn (2013) approach are the use of workers’ self-assessments, similar to Spitz-Oener (2006) and Autor & Handel (2013), and the level of analysis which is at the

[1] While the emphasis on codifiability is sensible, there is a danger of circularity: if routine tasks are defined as codifiable tasks that are being replaced by machines, technological change is by construction routine-biased.

[2] A third approach is to classify routine occupations based on census 1-digit occupational codes as in Acemoglu & Autor (2011). However, this very coarse method is clearly inferior to either expert-coded scores or worker self-assessments (Salvatori 2018).

'job'-level (2-digit occupations in 2-digit sectors). The measures developed by these authors have mainly been adopted in a European context (Eurofound 2014, 2017).

Table 5.1 contrasts both operationalisations of the RBTC theory. The overview documents the proliferation of empirical studies based on the approach developed by Autor et al. (2003) and Autor & Dorn (2013). Their methodology has been adapted to a range of contexts beyond the United States, including studies of individual European countries and comparative studies. Some authors, like Acemoglu & Autor (2011), Autor & Handel (2013), and Spitz-Oener (2006), only adopt the analytical framework but use different data while others, such as Goos et al. (2014) embrace the approach wholesale. The alternative approach of Fernández-Macías & Hurley (2017) has also been used in two Eurofound (2014, 2017) reports in which the same authors were involved. While other measures of occupational routine intensity do exist (e.g. Salvatori 2018), they do not generally make the same claim to generality as the Autor & Dorn (2013) and Fernández-Macías & Hurley (2017) measures. Therefore, my discussion focuses on these two prominent measures. The approach proposed in this chapter owes much to the work of Fernández-Macías & Hurley (2017) but goes beyond their important contribution in several dimensions.

Table 5.1: Comparison of the operationalisation of task content in key articles

	Autor, Levy & Murnane 2003	**Eurofound 2014/Fernandez-Macias & Hurley 2017**
Dimensions	routine manual; routine cognitive; non-routine interactive; non-routine analytic; non-routine manual	Routine; cognitive (also: social interaction; trade intensity)
Variables	**Routine manual:** finger dexterity. **Routine cognitive:** adaptability to situations requiring the precise attainment of set limits, tolerances and standards. **Non-routine interactive:** direction, control and planning. **Non-routine analytic:** quantitative reasoning requirements. **Non-routine manual:** eye-hand-foot coordination.	**Routine:** repetitive hand or arm movements; repetitive hand movements of less than 1 or 10 minutes; monotonous tasks; dealing with unforeseen problems. **Cognitive:** complex tasks; use of computers at work; use of internet at work; number of years of formal education necessary.
Data source	Dictionary of Occupational Titles (US; 1977 and 1991 versions)	European Working Conditions Survey (2010 wave)
Unit of analysis	3-digit census occupations (US)	"jobs": 2-digit occupation, 2-digit industry cells (EU-28)
Variants also used by	Spitz-Oener 2006 (for Germany, dimensions only); Goos & Manning 2007 (for UK); Acemoglu & Autor 2011 (also using O*NET); **Autor & Dorn 2013;** Autor & Handel 2013 (dimensions only); Goos, Manning & Salomons 2014 (for 16 EU countries); Mahutga, Curran & Roberts 2018 (for LIS countries).	Eurofound 2014, 2017

5.3 A critique of existing operationalisations of occupational task content

The thesis of this study is that the measures just described suffer from conceptual and empirical problems which can, however, be addressed. Conceptually, there are two major issues:

H 1.1: Existing approaches use ill-defined auxiliary task dimensions (secondary task dimensions other than routine intensity).

H 1.2: The variables used in existing approaches do not capture the concepts they are purportedly measuring.

Empirically, I identify three main shortcomings:

H 2.1: Existing approaches do not account for task variation within occupations by using expert-coded data.

H 2.2: Existing approaches fail to account for change over time within occupations.

H 2.3: Existing approaches do not account for differences between countries.

This section discusses these problems and how I propose to address them. I then introduce the new measures, analyse how they compare to existing indices and illustrate the arguments that I have discussed theoretically.

5.3.1 Ill-defined auxiliary task dimensions

The uniting feature of all approaches discussed above is their interest in occupational routine-intensity. However, all studies define one or more auxiliary axes of occupational tasks. For example, in Autor et al. (2003) that is the cognitive-manual axis and in Fernández-Macías & Hurley (2017) simply a cognitive task dimension. Caines et al. (2017) juxtapose routine and complex tasks. Usually, however, these auxiliary axes do not serve a well-defined purpose. Only Caines et al. (2017) formulate a theory of task complexity in relation to routine-intensity and technological change. In most other analyses, the auxiliary task dimension does not add much of substantive interest. For instance, Autor et al. (2003) do not posit any independent relationship between technological change and cognitive and manual task inputs. However, if the measure of routine-intensity is used to operationalise RBTC, whatever auxiliary measure is part of the analysis ideally should have some independent theoretical interpretation. In particular, since RBTC is often presented as an alternative to SBTC, it would be eminently helpful to have a measure for the latter that is constructed in a similar manner as the measure of RBTC.

5.3.2 Variables that do not capture key concepts

The variables used to operationalise cognitive and manual routine tasks in Autor et al. (2003) and Autor & Dorn (2013) completely fail to capture key aspects of the notion of routine as defined above, most importantly, repetitiveness. For example, they measure cognitive routine intensity with "adaptability to situations requiring the precise attainment of set limits, tolerances and standards", a criterion which appears geared towards low-level clerical jobs and jobs in the manufacturing sector (Autor et al. 2003, p. 1323). Manual routine intensity is measured with finger dexterity. The relationship between finger dexterity and codifiability and repetitiveness seems altogether questionable. It rests on the assumption that tasks which involve fine movements and coordination are repetitive and can be automated — a shaky assumption, for example, with regard to musicians and artisans who often require a great deal of finger dexterity. In Autor & Dorn (2013), the measure for routine tasks is the simple average of those two variables. Thus, even though repetitiveness is at the core of the concept, it barely features in the variables used to measure routine intensity. Overall, the variables used appear to have been chosen not so much with the abstract concept of routine in mind but rather with a preconceived set of purportedly routine occupations.

Like this study, Fernández-Macías & Hurley (2017) motivate their article with a desire to improve the match between concepts and operationalisations of task content and an explicit critique of elements of Autor et al.'s (2003) method. Their index comprises five items from the EWCS. Three questions on repetitive arm or hand movements and short repetitive tasks capture the repetitiveness dimension of routine, with the first identifying manual routine tasks and the second and third capturing repetitive tasks more broadly. A fourth question about monotonous tasks introduces the notion of 'boringness'. They also include a question on dealing with unforeseen problems which arguably is an inverted measure of codifiability — something that is unforeseen cannot be codified. However, above all, dealing with unforeseen problems requires creativity and problem-solving ability, two of the qualities measured by the complexity index. The Fernández-Macías & Hurley (2017) cognitive index focuses perhaps too much on computer usage with two out of four variables and abandons the task-based framework by including a worker characteristic (average education). Thus, while their measures undoubtedly get much closer to the core of the respective concepts, further improvements are necessary.

5.3.3 Expert-coded occupation-level data

The Autor & Dorn (2013) task measures are derived from the Dictionary of Occupational Titles (DOT), in which expert coders assign scores that characterise occupations in the United States. More recent studies often use the Occupational Information Network (O*NET) database which offers a wider range of indicators than the DOT but is also based on data from American workers and only provides aggregated data at the level of occupations (Acemoglu & Autor 2011, Caines et al. 2017). Survey data, by asking people what they actually do in their job, are conceptually closer to the idea of the task-based approach; furthermore, survey data can provide a sense of the variability of tasks within an occupation. Spitz-Oener (2006) moreover points out that experts tend to underestimate the true changes in task content. A drawback of the survey approach may be measurement error introduced by respondents understanding questions differently or having different reference points. Nevertheless, the best way to find out what people do at work seems to be to ask the workers themselves and to take the variability of their answers into account.

5.3.4 Failure to account for change within occupations

The data used by Autor & Dorn (2013), Fernández-Macías & Hurley (2017) and most subsequent studies make it impossible to account for within-occupation change over time.[3] However, changing job tasks are a crucial component of technological change, as numerous studies make clear. A case study by Fernandez (2001) details how job tasks changed at a plant that underwent modernisation. Spitz-Oener (2006) finds that in Germany, within-occupation changes account for most of the change in aggregate task requirements. It is therefore clear that as the prevalence of occupations changes, so does their nature. A failure to account for this would result in underestimating the impact of technological change on the labour market.

5.3.5 Failure to account for differences between countries

Furthermore, all existing comparative studies fail to account for potential differences in task content between countries. Although Eurofound (2014) rightly point out that job tasks should be relatively similar across developed countries, the possibility that some occupations differ between countries should not be dismissed. For example, in lagging developed economies limited access to computer capital may retard computerisation in routine cognitive occupations (Fonseca et al. 2018). Indeed, Eurofound (2014) find that there are differences across countries, albeit small, regarding the demand for routine and cognitive tasks in a job, relative to that country's task distribution. Yet, many studies use the measures of Autor et al. (2003) and Autor & Dorn (2013) outside of the US. Fernández-Macías & Hurley (2017), while they do not use data from a country that is not part of the analysis, still only calculate one measure of task intensity for all countries. Yet ideally, country-specific measures of task content should be used for more detailed analyses.

[3] Sebastian & Biagi (2018) provide an overview of the years from which task data in various studies are taken.

5.4 Towards better measures of task content

5.4.1 Meaningful auxiliary task dimension

Much of the RBTC literature lacks a well-defined auxiliary dimension to the routine dimension. This chapter proposes a measure with the aim of enabling an analysis of SBTC alongside RBTC. Following Caines et al. (2017), this is called the task complexity dimension and is defined as the demand for higher-order skills such as effective communication, abstraction and decision making. Occupations that comprise many tasks requiring these skills are less likely to be replaced by technology, as machines and artificial intelligence (AI) cannot (yet) perform such tasks. On the contrary, these higher-order skills are inclined to be complemented by modern technology: thanks to it, effective communicators can reach wider audiences, scientists have powerful tools at hand that facilitate abstraction and induction, and business leaders can make better decisions based on improved data. Hence, task complexity is suitable for measuring the prevalence of SBTC.

Routine task intensity (RTI) and task complexity are not just two sides of the same coin: some routine occupations also require performing a considerable number of complex tasks, for example some health and clerical occupations. Thus, task complexity is expected to be negatively correlated with routine intensity but is nevertheless analytically distinct.

5.4.2 Variables that capture key concepts

The operationalisation of RTI based on Autor et al. (2003) does not in fact measure RTI in the way they themselves define the concept: the key notions of repetitiveness and codifiability are insufficiently captured. Thus, there is a real need to better align the concept of routine-intensity and its measurement. Fernández-Macías & Hurley (2017) have taken an important step in this direction and this discussion of the issue relies heavily on their previous work, yet their measure of routine-intensity can be improved further.

My routine index includes the following items: whether a job involves (1) repetitive arm or hand movements, (2a) short repetitive tasks of less than 1 minute, (2b) short repetitive tasks of less than 10 minutes, (3) monotonous tasks and (4) meeting precise quality standards.[4] These five items are less occupation-specific than the ones used by Autor et al. (2003). At the same time, these items afford the notions of repetitiveness and codifiability their due importance. So, while this index departs almost completely from Autor et al. (2003), it closely resembles the RTI measure of Fernández-Macías & Hurley (2017).

The only difference to Fernández-Macías & Hurley (2017) is the inclusion of the item 'meeting precise quality standards' which is included instead of 'solving unforeseen problems on one's own'. Below, I show that the latter is better suited as a component of the complexity index. Fernández-Macías & Hurley (2017) and Eurofound (2014, p. 48) argue explicitly against the inclusion of a quality control variable, on the grounds that this assigns relatively high routine scores to higher-skilled occupations which often include monitoring tasks. They are correct, however, there is a crucial difference between enforcing quality standards and being forced to meet them. Therefore, the requirement to meet precise quality standards is a suitable indicator of codifiability in an index which might otherwise focus too heavily on the repetitiveness aspect of routine-intensity.

The task complexity index has no direct counterpart in Autor et al. (2003) and is also where I depart further from Fernández-Macías & Hurley (2017) who focus on cognitive intensity rather than task complexity as the second dimension of occupational task content. It aims to measure the demand for higher-order skills such as effective communication, abstraction and decision making. It includes the following items from the EWCS: whether a job entails (1) working with computers, tablets, smartphones, etc., (2) solving unforeseen problems on one's own, (3) complex tasks and (4) learning new things.[5]

The only overlapping question with Fernández-Macías and Hurley (2017) is whether a job involves complex tasks. The included questions reflect the fact that

[4]These items correspond to questions number 30e, 48a, 48b, 53a and 53d in the EWCS.
[5]These items correspond to questions number 30i, 53c, 53e and 53f in the EWCS.

on-the-job learning is a key characteristic of complex jobs, as is solving unforeseen problems on one's own.[6] Moreover, rather than two separate questions whether a job involves the use of computers and use of the internet, one question asking for the use of 'computers, tablets, smartphones, etc.' is included. This serves to avoid an undue emphasis on office jobs since it is unlikely that a job involves the use of computers but not the internet, or vice versa. Overall, using variables that truly capture the prevalence of routine and complex tasks in line with their definitions is an important and long overdue advance in the literature on RBTC.[7] It ensures that the indices measure what they claim to be measuring (criterion validity). The present indices are an important step in this direction.

5.4.3 Individual-level survey data

Using the EWCS helps to realise the advantages of individual-level survey data. Instead of one number assigned by an outsider, each score is the result of many (often thousands of) practitioners of an occupation evaluating what they do in their job and how they do it. It is worth noting that with O*NET replacing the DOT, there is a trend towards using survey data even in the US context. However, O*NET still only provides aggregated data at the occupational level (Acemoglu & Autor 2011). In contrast, a crucial benefit of using survey data is that it provides an indication of the degree of variation within a job (Autor & Handel 2013). Overall, the present approach follows the practice of a few scholars, mainly from Europe, who have used survey data for the analysis of occupational task content all along.

5.4.4 Accounting for change within occupations

With the EWCS, it is possible to analyse change within occupations in most European countries. In principle, any data source with consistent occupation-level

[6] Differences in education and training systems may affect the need for on-the-job learning in an otherwise similar occupation. However, the impact of this on overall complexity scores is likely to be small.

[7] Of course, no single survey question perfectly encapsulates complexity or routine intensity. This is why my indices comprise different items that capture different facets of complexity and routineness.

data for several points in time can be used for this purpose. However, the DOT was only updated infrequently and was eventually replaced by O*NET. It is therefore impossible to develop a time series of occupational task content based on the variables in Autor et al. (2003). Overall, for comparative analysts, the EWCS is the most appropriate available data source. Eurofound (2014) and Fernández-Macías & Hurley (2017) only use the 2010 wave and ignore the temporal component of occupational change. In contrast, I use data from the four most recent waves (2000 – 2015) to investigate change within occupations over a period of 15 years. Furthermore, the dataset can easily be expanded to include future waves of the EWCS once they become available.

5.4.5 Accounting for differences between countries

The EWCS data can also be used to analyse differences between countries. The EWCS covered 35 European countries in its most recent wave, with a target sample size between 1,000 and 4,000 individuals depending on country size.[8] With these characteristics, a country-level analysis at the 2-digit level for occupations is feasible within Europe. For comparative analyses including non-European countries, a pooled measure should be used, but compared to the prevailing practice of using task content measures based on US data for all countries, this broader approach promises more robust insights.

Table 5.2 provides an overview of how my measures differ from and improve upon exiting indices. Harking back to the five thesis statements above, it is clear that the proposal in this chapter entails important conceptual and empirical improvements over both Autor et al. (2003) and Fernández-Macías & Hurley (2017). With regard to the first overall research question of the thesis, it promises to better align the measurement of routine intensity with the underlying concepts of repetitiveness and codifiability, and account for relevant variation in task content over time and between countries.

[8]There are a few exceptions: the target sample size was 500 in Luxembourg in 2000 and 600 in Cyprus, Estonia, Luxembourg, Malta and Slovenia in 2005.

Table 5.2: Summary of the indices with their differences and improvements

Characteristics	Autor et al. 2003/Autor & Dorn 2013	Eurofound 2014/Fernandez-Macias & Hurley 2017	This proposal
Task dimensions	Routine (measure of predominance over abstract and manual tasks)	Routine, cognitive → measure of prevalence	Routine, complex → dimensions capturing RBTC and SBTC
Variables	Severe mismatch between concepts and variables	Remaining problems: mis-classification of 'dealing with unforeseen problems'; mixing of worker and task characteristics	Coherent operationalisation of routine and complexity dimensions
Data and unit of analysis	1977 DOT (O*NET in some later applications); US census occupations	2010 EWCS → contemporary individual-level survey data; ISCO-88 2-digit, NACE 2-digit 'jobs'	2000, 2005, 2010, 2015 EWCS → contemporary individual-level survey data, repeated cross-section; ISCO-88 2-digit occupations
Treatment of change over time	Limited analysis in Autor et al. (2003), none in other articles	No analysis of change over time	Detailed analysis of change over time
Treatment of country differences	Index specific to the US, no analysis of country differences	No analysis of country differences	Analysis of country differences in EU-15

5.5 Data and construction of the indices

As stated in chapter 4, the task measures are calculated based on waves 3 – 6 of the EWCS which provides workplace information for the years 2000, 2005, 2010 and 2015. The EWCS allows me to characterise the task profile of occupations with a set of relatively objective questions and a number of more subjective items. The precise set of questions comprising each index is as follows (summary statistics are in table 5.3 below).

Routine intensity index:

1. Does your main job involve repetitive hand or arm movements? (Scale from 1: “all of the time” to 7: “never”)

2a. Does your job involve short repetitive tasks of less than 1 minute? (Scale from 1: ”yes” to 2: “no”)

2b. Does your job involve short repetitive tasks of less than 10 minutes? (Scale from 1: ”yes” to 2: “no”)

3. Does your main paid job involve monotonous tasks? (Scale from 1: ”yes” to 2: “no”)

4. Does your main paid job involve meeting precise quality standards? (Scale from 1: ”yes” to 2: “no”)

Task complexity index:

1. Does your main paid job involve working with computers, laptops, smartphones etc.? (Scale from 1: “all of the time” to 7: “never”)

2. Does your main paid job involve complex tasks? (Scale from 1: ”yes” to 2: “no”)

3. Does your main paid job involve solving unforeseen problems on your own? (Scale from 1: ”yes” to 2: “no”)

4. Does your main paid job involve learning new things? (Scale from 1: ”yes” to 2: “no”)

I calculate three versions of the indices. The overall indices use all available data and are pooled across countries and waves. They are available for the EU-27 and EU-15 groups of countries. I predominantly use the EU-15 version in the comparative analyses throughout this thesis. In addition, I calculate a wave-specific version of the overall indices and a country-specific version with data for the respective country pooled over all available waves. Following the approach in most of the comparative literature, the indices are calculated at the 2-digit level for occupations and not for occupation-industry cells. This entails a certain loss of precision compared to Fernández-Macías & Hurley (2017) but in my view is warranted for two reasons. Firstly, only by restricting the analysis to the 26, 2-digit ISCO-88 occupations, is the analysis of country- and wave-specific scores feasible. Secondly, the method of Fernández-Macías & Hurley (2017), using a sample of about 43,000 to populate 3,123 out of 3,784 hypothetical job cells, means that for a large number of small jobs, the task scores are based on very few observations.[9] Some additional methodological details on the construction of the indices are reported in appendix B.

The indices are constructed by standardising the constituent variables to have a mean of 0 and a standard deviation of 1 following Acemoglu & Autor (2011), after reverse-coding the answers where appropriate. I then first average across individual survey respondents and subsequently across 2-digit ISCO codes. Principal component analysis, which is sometimes used in the literature, is not useful in the present case because of the low number of items that make up each index. Thus, the routine index *rscore2d* for occupation o is calculated as:

$$rscore2d_o = \frac{\sum_{i \in I_o} \left(\frac{\sum_{j \in J} ewcs_{ji}}{J} \right)}{I_o}, \tag{5.1}$$

[9]For the same reason, ISCO group 62 has been merged with group 61. This change affects 0.4 percent of observations in the original dataset.

where $ewcs_{ji}$ is the value for individual i on question j, J is the set of items used to calculate $rscore2d$, and I_o is the set of individuals with occupation o. Analogously, the complexity index $cscore2d_o$ is calculated as:

$$cscore2d_o = \frac{\sum_{i \in I_o} \left(\frac{\sum_{l \in L} ewcs_{li}}{L} \right)}{I_o}, \tag{5.2}$$

where L denotes the set of indicators used to calculate the index. For the wave- and country-specific indices, the respective subscripts w and c have to be added to the formula.

In table 5.3, I provide detailed summary statistics for the EU-27 version of the indices. The top two panels of the table show how respondents answered the individual questions, divided by sex and 1-digit occupational group. The percentages largely conform to intuitions about the character of occupations: for example, professionals have the lowest share of monotonous tasks and crafts and manufacturing workers exhibit the highest frequency of repetitive hand or arm movements. The bottom panel shows the means and standard deviations of the normalised indices. It shows no meaningful differences between women and men, and the scores for individual 1-digit occupational groups again are largely in line with theoretical expectations. Counterintuitive patterns appear to be due to the aggregation at the 1-digit level. For example, the marginally higher routine intensity of crafts compared to manufacturing occupations is a result of the classification of relatively low-routine drivers and mobile plant operators (group 83) as part of the manufacturing major group. Interestingly, the complexity index shows a greater degree of overall variation between occupational groups than the routine index.

Not accounting for sectoral differences may cause researchers to mistake differences in the sectoral composition of the economy for differences in occupational tasks and lead to a false positive finding of country differences in task content. Figure B.1 in appendix B details the extent of sectoral differences by major occupational group. It shows that the overall import of this is small: the outliers tend to be unlikely occupation-sector combinations such as clerks in agriculture or machine operators and assemblers in education which do not weigh heavily in the aggregate index. Goos & Manning (2007) for the UK, Acemoglu & Autor (2011) for the United States, and Fonseca et al. (2018) for Portugal likewise report that between-occupation variation far outstrips between-industry variation. Thus, since at present it is not feasible to address all potential sources of variation at once, I focus on the occupational level.

Table 5.3: EWCS task measures and indices by sex and ISCO-88 major group

	All	**Female**	**Male**	**Managers**	**Professionals**	**Assoc. professionals**	**Clerks**	**Service**	**Agriculture**	**Crafts**	**Manufacturing**	**Elementary**
Routine items												
Rep hand or arm												
Half or more	52.35	53.07	51.62	38.69	38.54	41.52	53.23	49.25	60.01	68.81	68.13	65.74
Less than half	47.65	46.93	48.38	61.31	61.46	58.48	46.77	50.75	39.99	31.19	31.87	34.26
Rep tasks <1min												
Yes	25.83	26.51	25.17	20.35	16.64	20.11	28.65	27.92	28.46	31.65	32.38	31.23
No	74.17	73.49	74.83	79.65	83.36	79.89	71.35	72.08	71.54	68.35	67.62	68.77
Rep tasks <10 min												
Yes	41.35	42.3	40.41	33.09	28.58	35.4	44.43	43.7	45.06	50.9	48.62	48.07
No	58.65	57.7	59.59	66.91	71.42	64.6	55.57	56.3	54.94	49.1	51.38	51.93
Quality standards												
Yes	71.31	68.2	74.4	71.3	71.93	72.68	67.34	65.1	63.18	87.58	77.07	61.53
No	28.69	31.8	25.6	28.7	28.07	27.32	32.66	34.9	36.82	12.42	22.93	38.47
Monotonous tasks												
Yes	44.16	44.7	43.62	36.77	30.25	33.69	47.35	44.67	53.94	48.96	57.99	59.16
No	55.84	55.3	56.38	63.23	69.75	66.31	52.65	55.33	46.06	51.04	42.01	40.84
Complexity items												
ICT use												
Half or more	37.96	40.67	35.28	55.75	60.74	60.25	75.06	21.45	3.71	11.76	14.19	7.90
Less than half	62.04	59.33	64.72	44.25	39.26	39.75	24.94	78.55	96.29	88.24	85.81	92.10
Unforeseen problems												
Yes	81.89	79.13	84.62	93.21	90.97	89.33	80.95	78.63	82.07	81.04	73.33	65.04
No	18.11	20.87	15.38	6.79	9.03	10.67	19.05	21.37	17.93	18.96	26.67	34.96
Complex tasks												
Yes	58.57	54.00	63.09	69.35	78.06	72.43	57.79	42.63	48.58	67.86	47.21	31.62
No	41.43	46.00	36.91	30.65	21.94	27.57	42.21	57.37	51.42	32.14	52.79	68.38
Learning new things												
Yes	70.12	69.27	70.97	77.60	91.03	85.86	73.21	62.06	54.55	70.82	54.57	40.18
No	29.88	30.73	29.03	22.40	8.97	14.14	26.79	37.94	45.45	29.18	45.43	59.82

Table 5.3 continued from previous page

	All	Female	Male	Mana gers	Profes- sionals	Assoc. profes- sionals	Clerks	Service	Agri- cul- ture	Crafts	Manufac- turing	Elemen- tary
Means and standard deviations of composite indices												
Routine intensity	0	-0.0017	0.0017	-0.1499	-0.2105	-0.1344	0.0201	-0.0182	0.0607	0.2312	0.2147	0.13
	-0.618	-0.6259	-0.61	-0.5765	-0.5577	-0.5814	-0.6179	-0.6309	-0.6003	-0.5794	-0.6147	-0.6285
Task complexity	0	-0.0301	0.0298	0.2628	0.4064	0.3316	0.207	-0.2346	-0.3257	-0.1001	-0.3315	-0.5846
	-0.6872	-0.7092	-0.6633	-0.5996	-0.4876	-0.5377	-0.6545	-0.6504	-0.6024	-0.5911	-0.6582	-0.6447
N	107,488	53,484	53,998	8,914	15,888	15,054	12,617	17,508	3,096	13,916	8,237	12,258

Table 5.4: Correlation coefficients of items constituting the RTI and complexity indices in EU-27 countries

	Rep. Arm/ hand	1 min tasks	10 min tasks	Monotonous	Standards	Computer	Complex	Unforeseen	New things
Repetitive arm/hand	1								
1 min repetitive tasks	0.30***	1							
10 min repetitive tasks	0.34***	0.51***	1						
Monotonous tasks	0.26***	0.22***	0.23***	1					
Meet quality standards	0.14***	0.10***	0.12***	0.09***	1				
Item-rest correlation	0.40***	0.45***	0.48***	0.30***	0.15***				
Computer use	-0.09***	-0.06***	-0.06***	-0.09***	0.06***	1			
Complex tasks	-0.05***	-0.02***	-0.00	-0.03***	0.15***	0.19***	1		
Unforeseen problems	-0.02***	-0.02***	0.01***	-0.02***	0.20***	0.27***	0.29***	1	
Learn new things	-0.06***	-0.03***	-0.00	-0.10***	0.20***	0.31***	0.30***	0.42***	1
Item-rest correlation						0.34***	0.46***	0.34***	0.49***

Note: Correlations in the EU-15 sample are virtually identical. * p <0.05, ** p <0.01, *** p <0.001

5.5.1 Internal validity of the indices

Before comparing them to other indices, it is crucial to establish the internal validity (construct validity) of my measures. To this end, I investigate the correlations of the individual items in my indices, item-rest correlations, and the specificity of 2-digit occupations. The findings discussed here pertain to the EU-27 indices; however, identical results are obtained looking at the EU-15 version instead.

The correlations of the components of the indices that are displayed in table 5.4 support the notion that the indices are internally consistent measures. Since the correlations are calculated at the individual level, they are not very high, but it is clearly visible that they form two clusters, corresponding to the indices. All components of the RTI and complexity indices are positively and significantly correlated with the other components of the index they are part of, with correlation coefficients usually in the range between 0.2 and 0.5. On the other hand, they tend to be weakly negatively correlated with the components of the other index.[10]

The only outlier is the quality standards item in the routine index which is positively correlated with all components of both indices. This suggests that there is some ambiguity which index the item belongs to; indeed, the correlations are slightly stronger with the items in the complexity index. Yet, for the conceptual reasons outlined above, I believe that the question belongs in the RTI index. For all other items, the smallest absolute value of the correlation with a component of the own index substantially exceeds the maximal absolute value of the correlation with a component of the other index. This indicates that although they are correlated, the two indices nevertheless measure different underlying concepts. The table also shows that the item “dealing with unforeseen problems”, which Fernández-Macías & Hurley (2017) include in their RTI index, in fact belongs in the complexity index, since it is strongly positively correlated with the items in that index while it is essentially uncorrelated with all RTI items except the quality standards item.

[10] A strong negative correlation would just mean that the items are the inverse of one another, hence we look for a low absolute correlation for reassurance that the indices do not measure two sides of the same coin.

Table 5.4 also contains an item-rest correlation test. For this test, an item is correlated with a version of the index comprising all items except the one in question. The only case where there is some ambiguity is again the quality standards question which only exhibits a relatively weak correlation of 0.15 with a version of the RTI index calculated without this question. Most other correlations are between 0.30 and 0.50 and thus provide further evidence for the internal validity of the indices. Table 5.4 thus provides support for the argument that my indices are internally consistent measures (construct validity).

A final possible concern is the specificity of 2-digit occupations and the associated within-occupation variation. Since 2-digit occupations comprise very different numbers of more detailed 3-digit and 4-digit occupations, the degree of task variation explained by the 2-digit occupational code could possibly be very different. While group 52 comprises two 3-digit categories and two 4-digit categories, group 82 covers nine 3-digit categories and thirty-seven 4-digit categories. Thus, one might expect the standard deviation to be higher in 2-digit occupations that consist of a higher number of detailed occupations. However, there is no meaningful relationship with regard to either RTI or complexity between the individual-level standard deviations and the number of 3-digit or 4-digit groups that make up a 2-digit occupation. The correlation between the standard deviation of the RTI index and the number of 4-digit occupations is -0.16 – and thus in the opposite direction as the initial worry –, all other relationships are weaker still. Thus, there is strong support for my measures from the perspective of internal validity.

5.6 Comparing measures of task content

This section presents empirical evidence with regard to the first part of my overall research question 1. It describes the novel measures and compares them to the 'competitor' measures of Autor & Dorn (2013) and Fernández-Macías & Hurley (2017). By comparing them to established indices, I establish the criterion validity of my measures by showing that they measure what they purport to be measuring. Furthermore, this section shows how the novel measures address the concerns

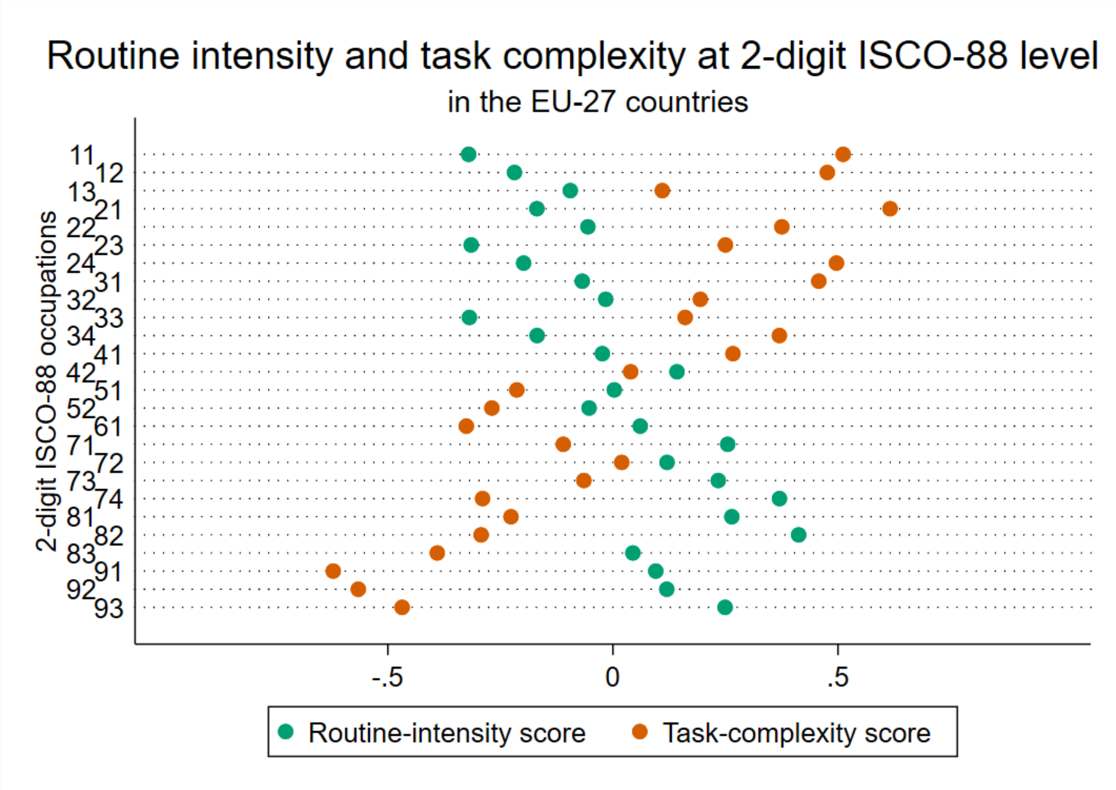

Figure 5.1: *Routine intensity and task complexity at the 2-digit ISCO-88 level.*

formulated in hypotheses 1.1, 1.2, and 2.1 and yield more plausible results for individual occupations. Thus, this section shows that occupational task content can be measured in a way that better corresponds to the underlying theoretical concepts.

5.6.1 Describing the RTI and complexity indices

A plot of the routine and complexity indices for the EU-27 at the 2-digit ISCO-88 level in figure 5.1 shows a relatively linear increase of routine-intensity down the occupational hierarchy and a countervailing decrease of task complexity. Thus, at least based on this ordering of occupations, routine occupations do not cluster around the middle of the occupational distribution.

Echoing the findings of previous research on routine tasks, cognitive tasks and complexity, the analyses find an inverse relationship between RTI and complexity. The Spearman correlation between the two indices is -0.73 and the weighted Pearson correlation, at -0.66, is in the range reported in other studies for the correlation between RTI and cognitive task intensity. This shows that it remains a challenge

Table 5.5: One-way analysis of variance

Response variable	SS	df	F (sig.)	ICC	Est. SD
Routine					
Between occupations	3,464.5	25	399.19 (0.00)	0.0888	0.1839
Within occupations	37,305.3	107,462			0.5892
Total	40,769.8	107,487			
Complexity					
Between occupations	12,685.7	25	1,473.59 (0.00)	0.2649	0.3522
Within occupations	37,004.1	107,462			0.5868
Total	49,689.8	107,487			

Note: The ratio of the estimated standard deviation over the combined estimated standard deviations gives the share of variation between and within occupations. SS: Sum of squares; df: Degrees of freedom; ICC: Intraclass correlation; Est. SD: Estimated standard deviation.

to develop an index based on a priori considerations with dimensions that do not explain the same underlying variation.

Table 5.5 reports analysis of variance (ANOVA) analyses which suggest that about 24 percent of the variation in RTI is between occupations, while 76 percent is between workers within occupations. For task complexity, the numbers are 37 percent between- and 63 percent within-variation. The corresponding intraclass correlations confirm that within occupations, individual observations are more similar when it comes to task complexity: the intraclass correlation is 0.26 for the complexity indices, compared to 0.09 for the RTI indices. The finding of a large component of within-occupation variation is in line with previous research (e.g. Spitz-Oener 2006) and illustrates the benefits of using survey data rather than expert-coded measures. While broad occupational categories play an important role in explaining what people do in their jobs, they cannot capture the full complexity of individual workplaces — especially with regard to the extent of variation in routine tasks required from people in comparable occupations.[11]

[11]The logic of the ISCO approach to ordering occupations is closely related to the logic behind the complexity dimension, as it is based on similarity in the skill level and skill specialisation of the tasks that make up a job. Routine intensity, on the other hand, is less directly linked to skills.

5.6.2 Comparing routine indices

The RTI measure of Autor & Dorn (2013), based on Autor et al. (2003), has become the standard 'off-the-shelf' measure for routine task intensity in the US and beyond. Fernández-Macías & Hurley (2017) improve on their measure with an approach which is more closely related to the method proposed here but requires further refinements. Therefore, comparing the new measure with these two RTI indices is paramount.

While Autor & Dorn (2013) work at the level of US census occupations, Goos et al. (2014) take their measure and map it onto ISCO-88, thus making the index applicable outside the US. The differences between Autor and Dorn (2013) and my RTI measure become visible in figure 5.2. The markers are dispersed widely over the plot region and the line of best fit from a weighted linear regression meets the y-axis nowhere near the origin. This, and the relatively low adjusted R^2 of 0.35, implies that my operationalisation is substantively different from the approach commonly adopted in labour economics.[12]

A look at the outliers is instructive to assess why the indices differ so much. Office clerks (group 41) have the highest RTI score of all occupations in Autor & Dorn (2013) but are just below the median according to my measure. This illustrates, in my view, the unrealistic characterisation of clerical occupations as far surpassing any other occupation in routine intensity. While secretaries, finance clerks, or librarians undoubtedly perform a fair share of routine tasks, it seems implausible that their job tasks are vastly more routine-intensive than those of printing machine operators, mechanical equipment assemblers, or weavers. Time likely plays a role here, since by 2015 clerical occupations undoubtedly had become less routine-intensive compared to 1977, the year from which the task data in Autor & Dorn (2013) are taken.

Overall, my measure tends to assign relatively higher routine-intensity scores to occupations in major groups 7, 8 and 9 (plant and machine operators, and assemblers; craft and related trades occupations; and elementary occupations). At the same time, managers, professionals and clerks tend to receive lower RTI

[12]From this section onward, the measures based on the EU-15 sample are used to maximise the overlap with the samples in chapters 6 and 7. Figures B.2, B.3, and B.4 in appendix B confirm that the relationship with existing measures is virtually identical in the EU-27 sample.

scores than in Autor & Dorn (2013). Most of this accords with the classical RBTC hypothesis, but the finding that elementary occupations — which include many low-skilled service jobs — are relatively high-routine contradicts the notion that displaced routine workers would move into such occupations that are classified as manual-interactive in Autor & Dorn (2013).

Next is the comparison with the measures proposed by Fernández-Macías & Hurley (2017). I use data from their Table 1 which provides aggregated data at the 2-digit ISCO-88 level that can be directly compared to my index (Fernández-Macías & Hurley 2017, p. 574). Here, the occupations align much better; discrepancies are mainly visible with regard to medium-to-high routine occupations. Nevertheless, my findings qualify their stance somewhat. Salespersons and elementary occupations all are less routine-intensive according to my measure. All three elementary occupations are among the five most routine-intensive occupations according to Fernández-Macías & Hurley (2017) , while only one (labourers in mining, manufacturing, construction and transport) makes the top-five according to my method. Conversely, blue-collar occupations in major groups 7 and 8 are closer to the upper end of the RTI scale on my index. It seems plausible that crafts and manufacturing occupations which frequently require a significant degree of job-specific training would be more standardised and repetitive than unskilled elementary occupations in which workers often take over low value-added tasks from their more skilled colleagues.

Overall, this suggests that Fernández-Macías & Hurley go too far when they claim that routine tasks are most frequent at the bottom of the "skills-wage-cognitive tasks continuum" (2017, p. 575). They underestimate the routine-intensity of crafts and manufacturing occupations, which are middling occupations in terms of wages and skills. The less skilled and lower-paid elementary occupations arguably comprise a less standardised set of tasks and are consequently slightly less routine-intensive. My index reflects this.

Thus, the three routine measures have different implications for occupational hierarchies. All are unanimous that managerial and professional occupations are the least routine-intensive; however, differences are visible when considering medium-

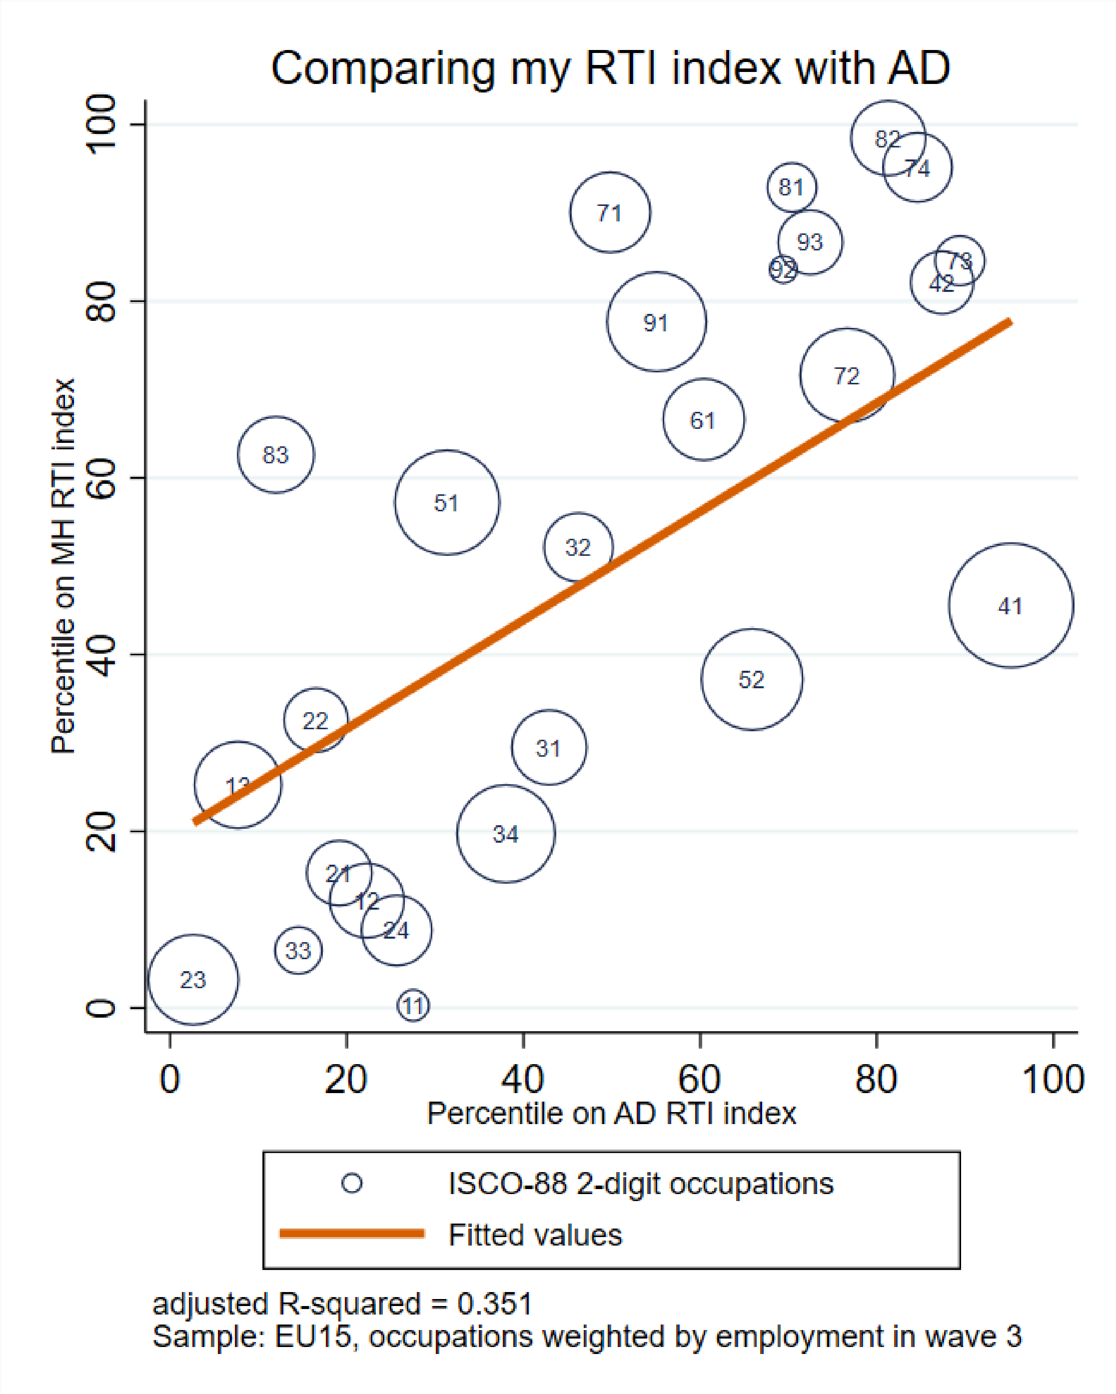

Figure 5.2: *Comparison of my routine index with Autor and Dorn 2013.*

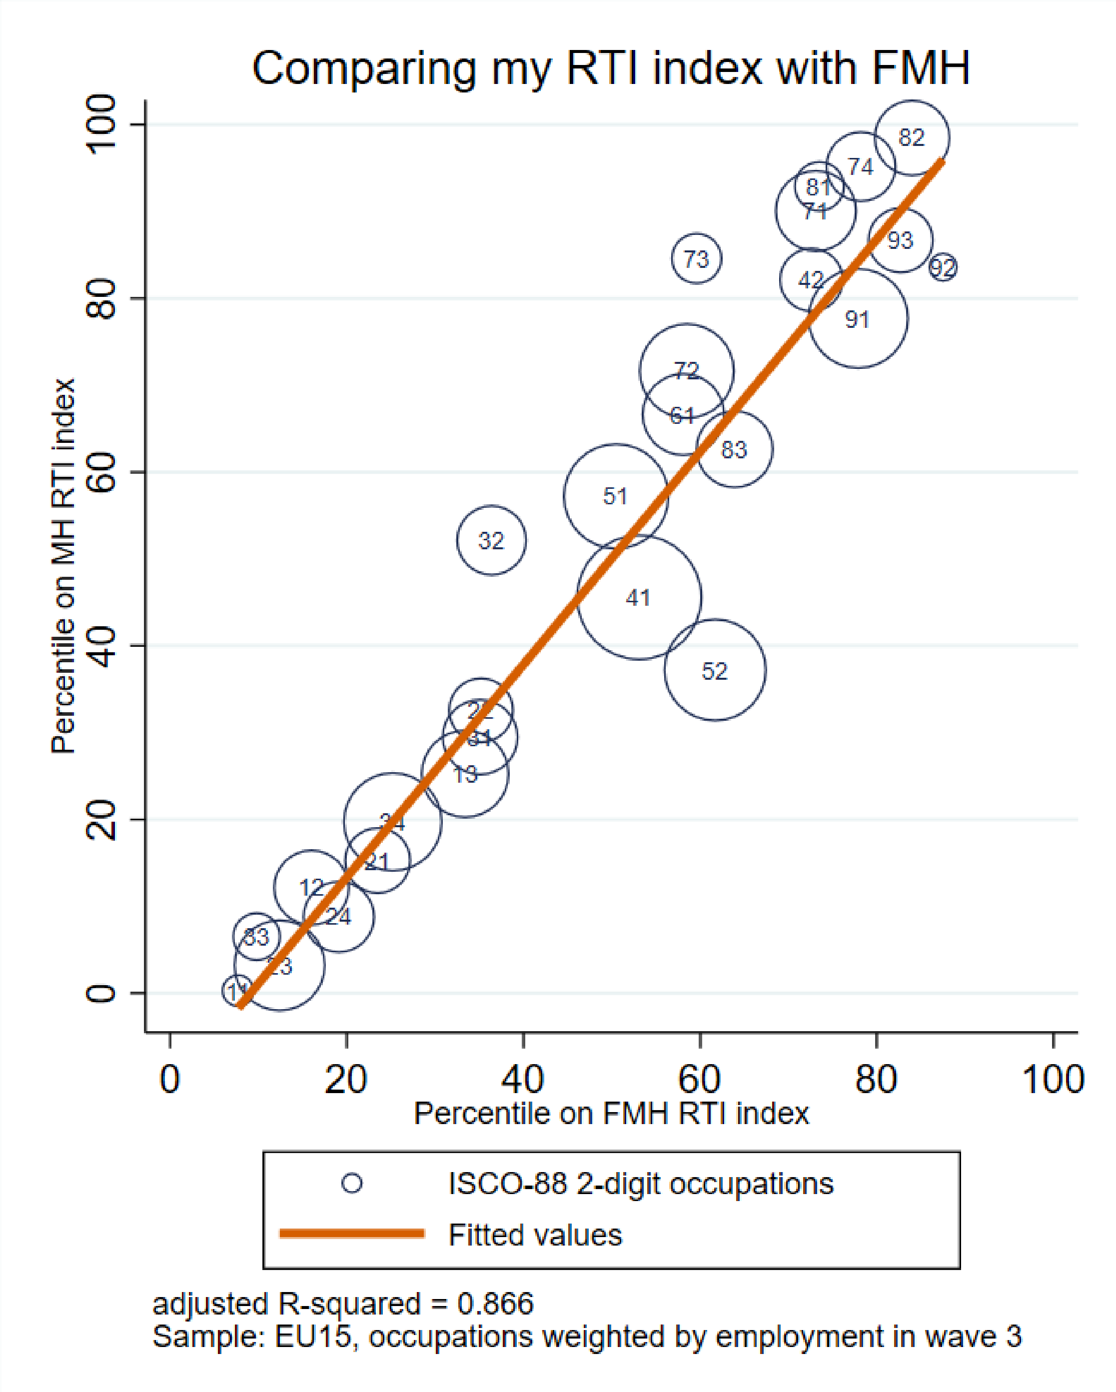

Figure 5.3: *Comparison of my routine index with Fernández-Macías and Hurley 2017.*

and high-routine occupations. Regarding clerical occupations and some crafts and manufacturing occupations, as well as service and elementary occupations, there are large discrepancies. Some of these occupations account for substantial portions of total employment. This shows that better alignment of the concept of routine-intensity and its measurement, compared to Autor & Dorn (2013) and Fernández-Macías & Hurley (2017), results in a partial reshuffling of the list of high-routine occupations. This, in turn, may have far-reaching consequences for studies of RBTC and employment and wage changes that build upon these different indices.

5.6.3 Comparing cognitive and complexity indices

My complexity index provides a measure of skill bias as discussed in section 5.4.1. It has its counterpart in the cognitive index that is proposed by Fernández-Macías & Hurley (2017) and in Eurofound publications involving the same authors (Eurofound 2014, 2017). However, they do not frame their index as a tool for analysing skill-bias but as the 'other side of the same coin' as their routine measure. Recalling section 5.5, the only direct overlap is that both measures ask whether a person's job involves complex tasks. Nevertheless, the resulting ordering of occupations is remarkably similar, as figure 5.4 shows.

Indeed, despite substantial differences in the construction of the two indices, the only larger discrepancies are between life science and health associate professionals, and teaching professionals and agricultural occupations. However, there is no broader group of occupations that is systematically ranked differently on the complexity index compared to the cognitive index. Thus, even though the concept and operationalisation are different, the practical implications of operating with task complexity rather than cognitive intensity are likely to be small.

To summarise the comparison, table 5.6 displays the rank order correlation between the various indices discussed and shows that the ordinal rankings of occupations are fairly similar. While some degree of similarity is to be expected, in the case of the complexity and cognitive indices it is striking how very different questions yield an almost identical ordering of occupations. The Autor & Dorn

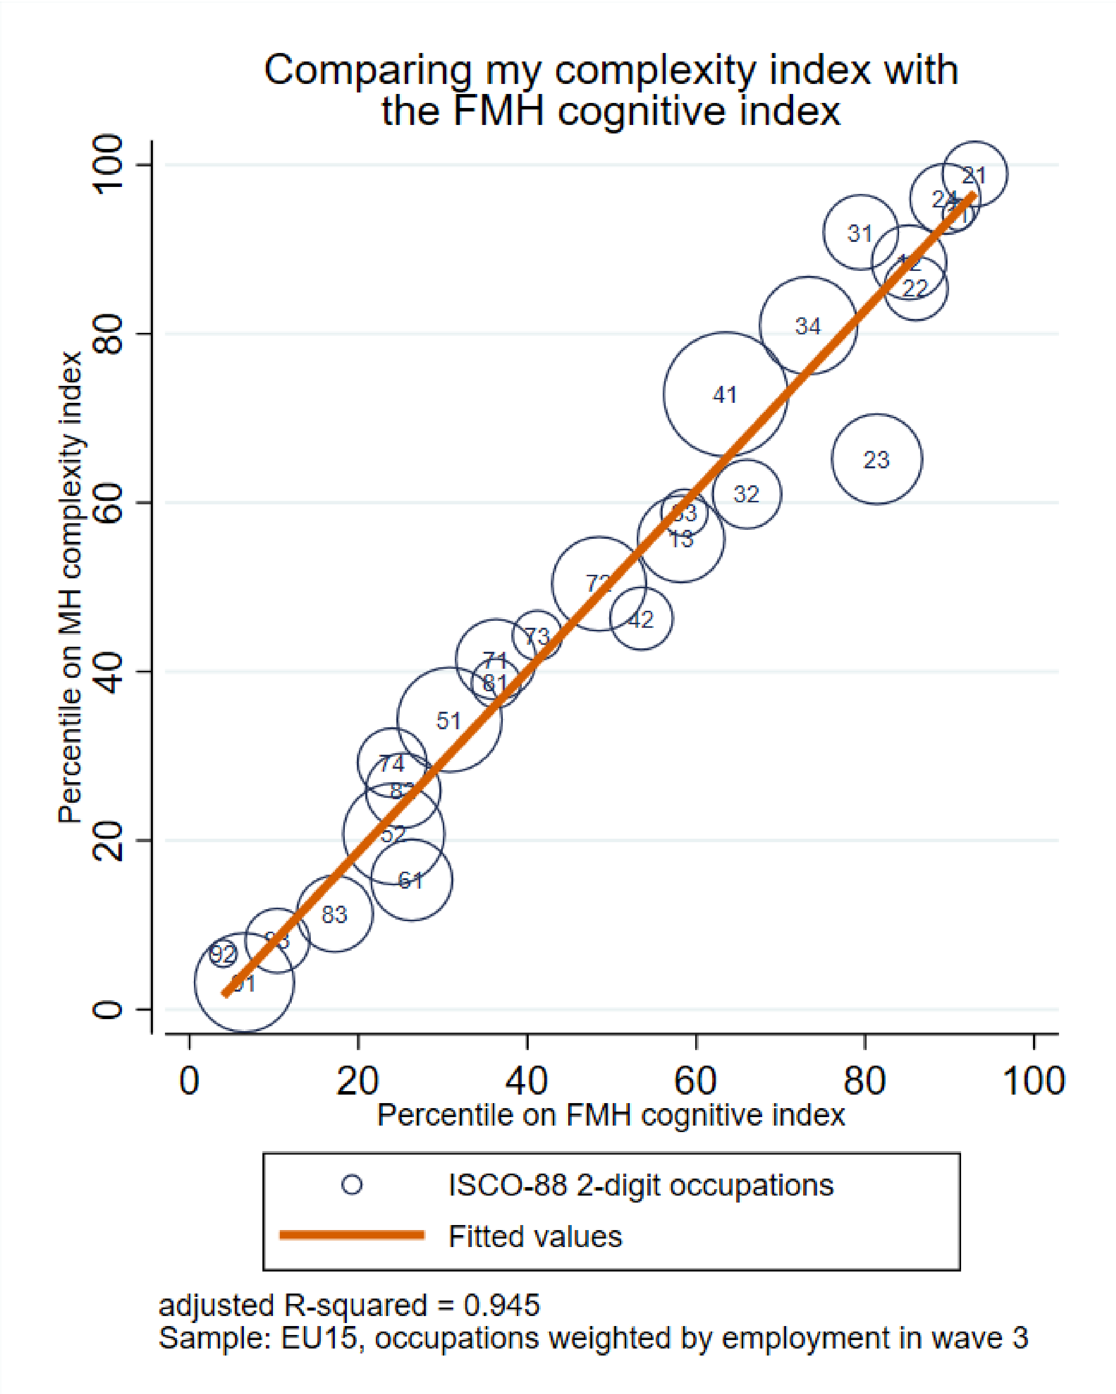

Figure 5.4: *Comparison of my complexity index with Fernández-Macías and Hurley 2017.*

Table 5.6: Rank order correlations between the indices

Index	AD RTI	FMH RTI	Complexity	FMH cognitive
RTI (this study)	.74	.94	-.73	-.76
AD RTI		.69	-.40	-.47
FMH RTI			-.83	-.86
Complexity (this study)				.98

Note: Calculated based on employment shares in the EU-15 in 2000 for my indices and the Autor and Dorn (2013, AD) RTI index, and taken from Table 1 in Fernández-Macías and Hurley (2017, FMH) for their RTI and cognitive indices.

(2013) index, calculated as a measure of predominance of routine tasks, exhibits comparatively lower correlations with my and Fernández-Macías and Hurley's (2017) routine indices which measure the prevalence of the respective tasks. The positive correlations with indices that attempt to measure the same underlying concept indicate that my indices exhibit criterion validity and in fact measure routine intensity and complexity – indeed, the theoretical discussion above suggests that they do so better than their competitor indices.

Furthermore, there are relatively strong negative correlations between the routine indices and the auxiliary measures of complexity and cognitive intensity. This is expected, yet it does not imply that they are two sides of the same coin, as there are several high-routine complex occupations and vice versa. Overall, the comparisons show that the routine intensity and task complexity indices do not simply replicate previous research. The complexity index is more conceptually meaningful, and the more theoretically informed operationalisation of both measures leads to a reappraisal of the task content of some occupations. Furthermore, the benefits of using survey data to analyse within-occupation variation are confirmed.

This validates the thesis statements 1.1, 1.2, and 2.1 that were formulated in chapter 5.3. Overall, the pooled indices of occupational task content based on the EWCS have been shown to capture the underlying concepts better than existing operationalisations, thus answering the call of research question 1 for better alignment of theory and measurement of occupational tasks.

5.7 Accounting for relevant variation: New opportunities for research

In addition to the improvements highlighted in the previous section, the measures developed in this chapter rely on superior data and so create exciting new opportunities for research into task change over time and differences between countries, thus addressing the concerns formulated in hypotheses 2.2 and 2.3. The analyses in this section furthermore speak to the second part of research question 1, as they show relevant variation over time and between countries that has been ignored by previous research.

5.7.1 Task content over time

It is widely accepted that not only the prevalence of occupations changes in response to technological advances, but also the tasks they entail (Spitz-Oener 2006). Yet, in the literature on RBTC, this facet of technology has received scant attention. The few studies that do look at within-occupation change over time find contradictory results. My index makes it possible to analyse within-occupation changes in the EU-15 countries over four points in time over a 15-year period from 2000 until 2015, thus covering a wider geographical area and a more recent time period than other studies. I find that there was no consistent trend towards less routine-intensive work, but a noticeable increase in overall occupational complexity. At the level of individual occupations, wave-on-wave increases in complexity are associated with a decline in routine-intensity.

Figure 5.5 illustrates the different patterns at the level of individual occupations. For RTI, there is no clear movement in one direction; if anything, low-routine occupations appear to have become marginally more routine-intensive. The 15-year differences are not statistically significant in most cases, however. By contrast, there is an almost uniform increase in task complexity, with statistically significant increases in 19 of the 26 occupations. The increases in task complexity tend to be larger in the more complex occupations, giving rise to further divergence between simple and complex occupations.

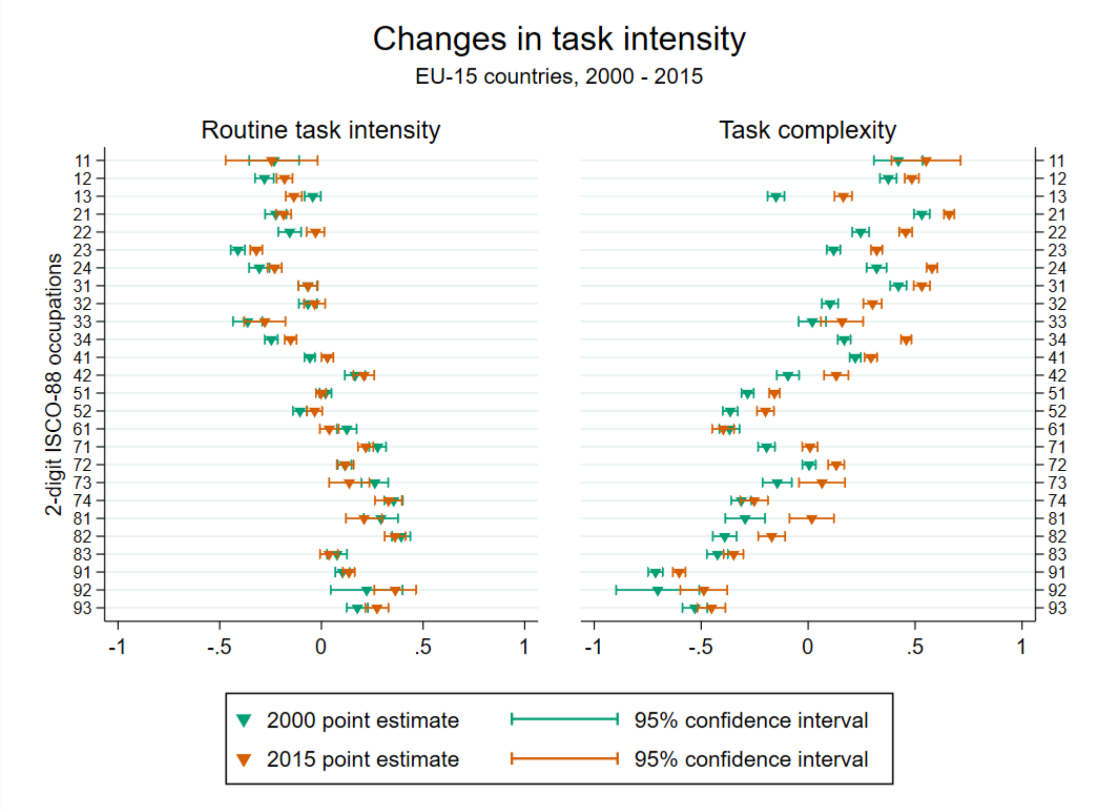

Figure 5.5: *Changes in average task intensity in the EU-15 countries from 2000 to 2015.*

These occupation-level takeaways are reinforced if one considers broad occupational groups, as can be seen in table 5.7. Occupations with above- and below-median routine-intensity converged in terms of RTI between 2000 and 2015, although this was not a linear process. These within-occupation changes constitute the intensive margin of task change. Over the same time, the combined employment share of the routine-intensive occupations declined from 51.3 percent to 48.2 percent, exemplifying the extensive margin of task change. Contrary trends at the extensive and intensive margins therefore result in trendless fluctuation of overall RTI.[13]

[13]Intuitively, low-routine occupations, which tend to be high-value-added occupations where potential gains from eliminating routine tasks would be high, would be expected to further reduce their routine task component. However, since computerisation and automation are not new phenomena, one possibility is that routine tasks have first been eliminated in high-value-added occupations but that this process had been largely completed by 2000 (the findings of Beaudry et al. 2016 about the reversal in the demand for skills and cognitive tasks may be interpreted in this vein). In the absence of further innovations, this may have led to routine tasks creeping back in, for example due to tighter regulations at the workplace. With each wave, a higher percentage of respondents gave a positive answer to the question, “does your main job involve meeting precise quality standards?” Hence, this may indeed explain part of the increase in routine task intensity through a regulatory channel in high-value-added occupations. Firm-level research over a longer

Table 5.7: Trajectory of the RTI and complexity measures, 2000 - 2015

Measure/occupations	**2000**	**2005**	**2010**	**2015**	**Δ**
Routine task intensity					
All occupations	-.043	-.078	-.034	-.002	+.041
Non-routine occupations	-.185	-.160	-.107	-.129	+.056
Employment share in %	*(48.7)*	*(48.2)*	*(50.0)*	*(51.8)*	*(+3.1)*
Routine occupations	.146	.105	.164	.108	-.038
Employment share in %	*(51.3)*	*(51.8)*	*(50.0)*	*(48.2)*	*(-3.1)*
Task complexity					
All occupations	-.094	.063	.111	.131	+.225
Simple occupations	-.369	-.261	-.293	-.237	+.132
Employment share in %	*(49.8)*	*(51.0)*	*(48.1)*	*(47.1)*	*(-2.7)*
Complex occupations	.203	.310	.316	.398	+.195
Employment share in %	*(50.2)*	*(49.0)*	*(51.9)*	*(52.9)*	*(+2.7)*

Note: Data for EU-15, occupations weighted by their employment share for the respective wave. Low-routine and simple occupations score below the median, high-routine and complex occupations score at or above the median of the respective measure.

Concerning complexity, both margins have reinforced each other and contributed to strong overall upskilling. Both simple and complex occupations are more complex in 2015 than they were in 2000, yet, they have diverged further. While simple occupations have become more complex by 0.13 points, complex occupations have increased in complexity by 0.2 points. That the overall increase is even bigger, by 0.23 points, highlights the contribution of compositional changes to changes in task complexity, as the employment share of complex occupations increased from 50.2 percent to 52.9 percent.

Finally, I consider to what extent changes in RTI and task complexity, are related. Plotting the wave-on-wave changes in RTI and task complexity for each occupation and regressing one on the other, a moderate negative relationship between changes in the two task dimensions is visible and depicted in figure 5.6. More precisely, a reduction of the RTI measure by 0.1 points is associated with an increase in complexity by roughly 0.3 points. Of course, there is no reason to assume that changes in complexity cause changes in routine-intensity or vice versa;

period would be necessary to determine the veracity of this conjecture, especially in light of contrary findings for the US (Hershbein & Kahn 2018).

rather, it stands to reason that an omitted variable – technological change – affects both simultaneously. Thus, even though there have been conflicting trends in low-routine and high-complexity occupations, broadly speaking, occupations that have become more complex have also become less routine-intensive, and vice versa. This relationship is robust to controlling for the survey wave which suggests that there has been no tendency for task change to accelerate or slow down in a linear fashion.

However, as even high-skill jobs become increasingly automatable (Acemoglu & Restrepo 2018), it is conceivable that in the foreseeable future more pronounced within-occupation changes will take place. In any case, this analysis has shown with cross-country data what so far has only been established empirically in a few individual countries: occupational task content has changed systematically over time. Furthermore, the intensive and extensive margin of change both contributed to the overall trends of broadly stable routine intensity and increasing task complexity.

5.7.2 Task content across countries

My proposal is also the first to account for country differences by calculating country-specific versions of the routine and complexity indices. The appeal of this is obvious: countries differ with regard to their sectoral composition, technological advancement, and labour relations, all of which may influence the prevalence of routine tasks in a given occupation. I find that there are indeed non-negligible differences between countries in the ordering of occupations, validating the argument that researchers should avoid applying one country's task data to another country, wholesale. Furthermore, a pooled index like the one I discussed in section 5.6 is shown to be a viable alternative when country-specific measures are not available for all countries in a sample.

The boxplots in figure 5.7 visualise the range of country-specific task intensity scores in the EU-15 countries. For most occupations, the country-specific RTI scores are in a narrow range around the median. In a number of small occupations, there is wider variation, but this is likely due to small sample sizes. The complexity index, in line with its greater overall variation (see section 5.6.1), shows greater

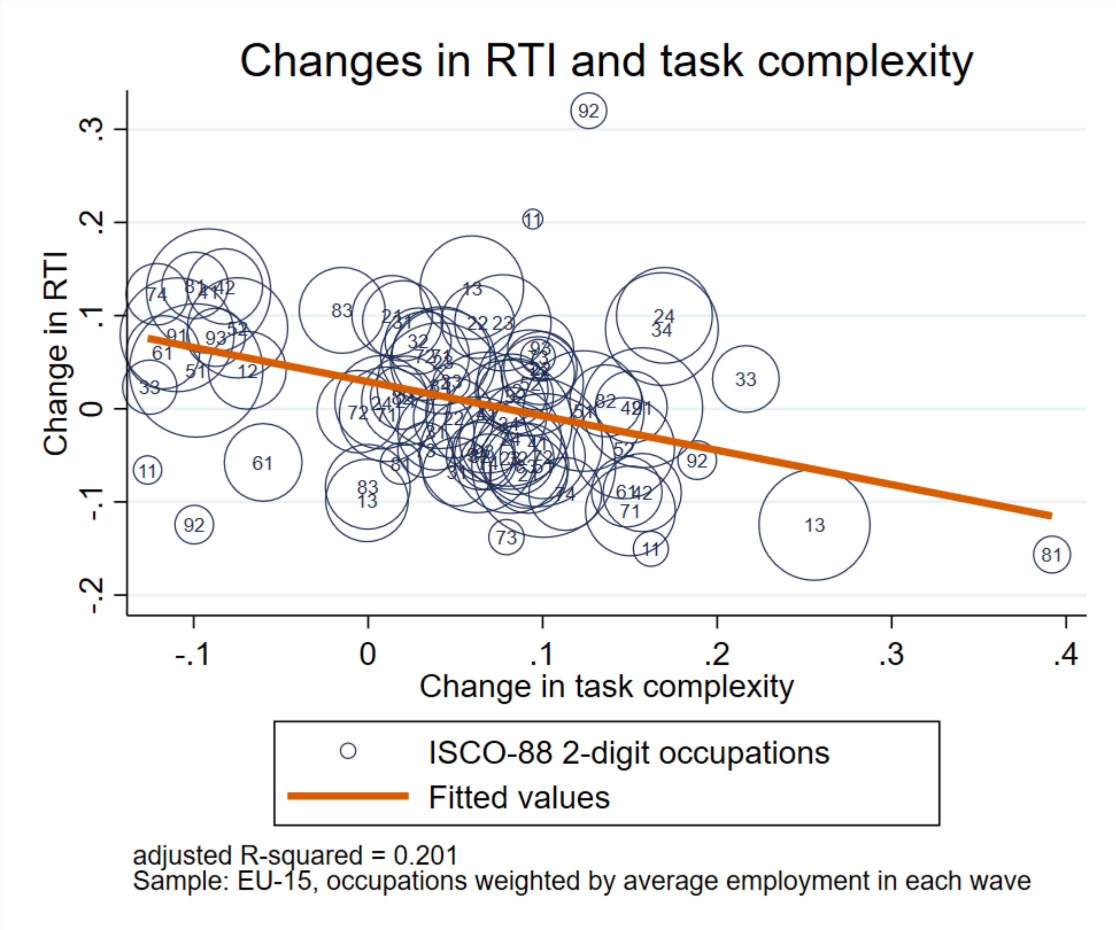

Figure 5.6: *The relationship between changes in RTI and complexity.*

between-country differences as well. This suggests that country-specific data are to be preferred unless small sample sizes make measurement imprecise.

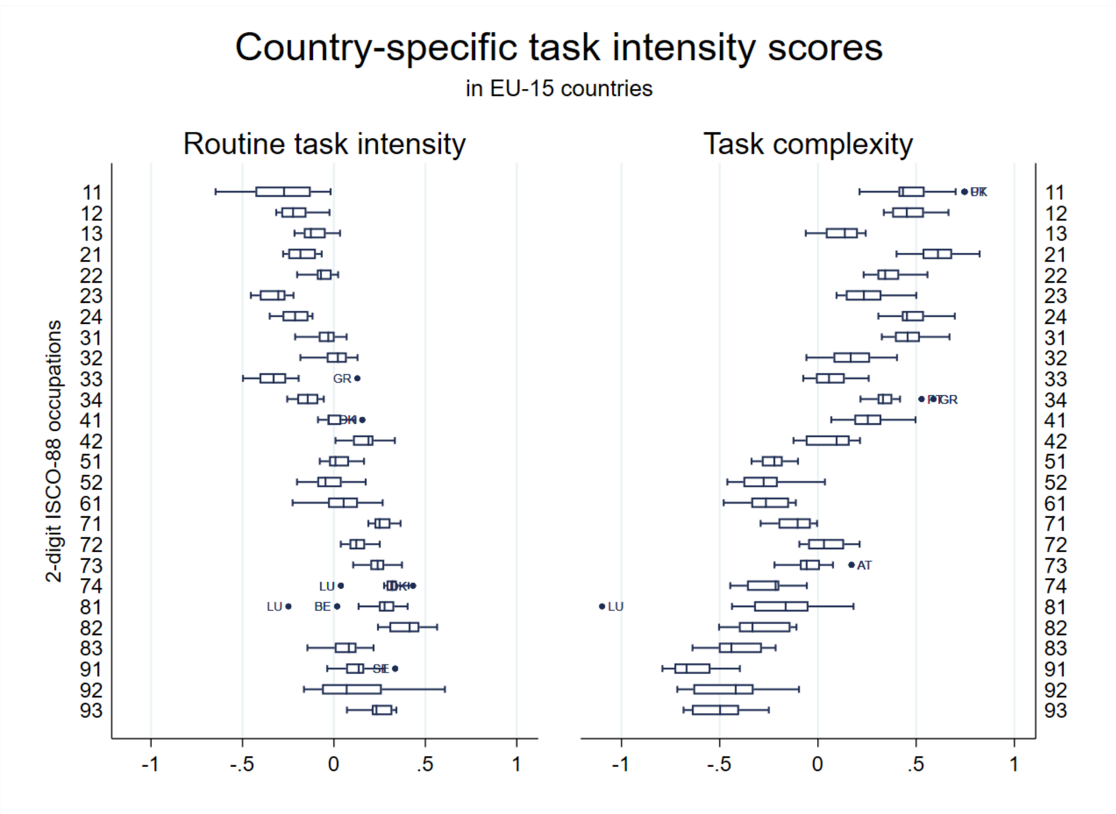

Figure 5.7: *Range of country-specific indices in the EU-15 countries.*

Table 5.8: Rank order correlation of the country-specific RTI and complexity rankings

	AT	BE	DE	DK	ES	FI	FR	GR	IE	IT	LU	NL	PT	SE	UK	EU-15
AT		0.96	0.94	0.96	0.92	0.94	0.93	0.92	0.95	0.93	0.91	0.97	0.94	0.95	0.96	0.97
BE	0.74		0.95	0.94	0.96	0.92	0.97	0.94	0.95	0.95	0.92	0.96	0.95	0.96	0.95	0.98
DE	0.84	0.87		0.95	0.91	0.92	0.94	0.96	0.94	0.94	0.94	0.92	0.90	0.94	0.93	0.96
DK	0.80	0.83	0.87		0.92	0.92	0.93	0.92	0.97	0.95	0.93	0.94	0.95	0.97	0.95	0.97
ES	0.95	0.75	0.84	0.84		0.86	0.94	0.77	0.89	0.92	0.80	0.95	0.95	0.84	0.81	0.97
FI	0.83	0.86	0.86	0.90	0.79		0.86	0.74	0.86	0.90	0.86	0.90	0.90	0.85	0.79	0.92
FR	0.82	0.82	0.82	0.81	0.85	0.87		0.94	0.96	0.92	0.86	0.95	0.92	0.92	0.95	0.97
GR	0.68	0.66	0.72	0.72	0.66	0.82	0.75		0.89	0.85	0.87	0.90	0.78	0.90	0.89	0.81
IE	0.76	0.82	0.85	0.84	0.77	0.84	0.83	0.71		0.94	0.89	0.95	0.89	0.94	0.94	0.93
IT	0.92	0.82	0.89	0.90	0.95	0.79	0.89	0.68	0.84		0.89	0.92	0.94	0.92	0.86	0.95
LU	0.77	0.82	0.73	0.69	0.71	0.60	0.77	0.41	0.65	0.73		0.88	0.81	0.95	0.90	0.83
NL	0.82	0.79	0.83	0.91	0.84	0.85	0.89	0.69	0.76	0.89	0.76		0.93	0.95	0.95	0.97
PT	0.73	0.65	0.83	0.66	0.64	0.69	0.57	0.63	0.74	0.65	0.51	0.57		0.85	0.80	0.97
SE	0.74	0.87	0.85	0.89	0.69	0.91	0.80	0.74	0.81	0.72	0.61	0.86	0.69		0.95	0.87
UK	0.81	0.86	0.86	0.88	0.76	0.88	0.79	0.65	0.94	0.83	0.68	0.79	0.79	0.86		0.85
EU-15	0.90	0.86	0.90	0.92	0.94	0.89	0.93	0.75	0.88	0.96	0.74	0.92	0.70	0.84	0.88	

Note: upper triangle: complexity index – lower triangle: RTI index.

More importantly, however, I consider the differences in the ordinal ranking of occupations across countries. The analysis shows that there are indeed strong reasons to prefer country-specific or pooled data for cross-country analyses. Tabulating the rank order correlations of the country-specific measures shows that the ordering of occupations in terms of routine-intensity and task complexity is by no means identical even in countries as similar as the EU-15. The correlations are displayed in table 5.8, with the lower triangle containing the correlation between RTI rankings and the upper triangle that of the complexity indices. It shows that the average correlation between two countries' RTI ranking is of a similar magnitude as the correlation between my pooled RTI index and the Autor & Dorn (2013) index, at 0.78. At the same time, the average correlation of each individual country's RTI ranking with the pooled index is substantially higher, at 0.87. Only 18 percent of the 105 country dyads show a higher correlation, and not a single country has a higher average correlation with the 14 other countries in the sample, driving home the point that it is highly problematic to assume that occupational routine intensity is constant across countries.

The differences are less pronounced with regard to the complexity index, with an average correlation of 0.93 between countries and the pooled index, and 0.91 between country dyads, of which a full 52 percent have a correlation greater than 0.93. Overall, because of the clear advantages associated with the country-specific measures of routine-intensity, analysts should use such measures wherever that is feasible. If it is not, this analysis shows that an index pooled over many countries will, on average, still be substantially closer to a country's true ranking than an index based on data from a single country, at least when it comes to RTI. The pooled RTI index presented here therefore constitutes a viable alternative to existing measures in cross-country analyses that extend beyond the coverage of the EWCS.

Overall, this section shows that differences in task content over time and between countries deserve greater attention. Hypotheses 2.2 and 2.3 thus receive empirical support. My indices provide the tools for analysing such differences in the EU-15 countries. Together with the factors discussed in section 5.6, this constitutes a significant improvement not only over the indices of Autor & Dorn (2013) but also Fernández-Macías & Hurley (2017), and comprehensively addresses the first research question of this thesis.

5.8 Conclusion

The literature on occupational task content has long relied on just a few off-the-shelf measures without giving much thought to how key concepts are theorised and operationalised. Hence it came to be that cognitive tasks were defined without a clear purpose, routine tasks operationalised with unsuitable variables and task variation within occupations, ignored. Further ignored were change over time and diversity across countries. Yet, for a meaningful comparative analysis of the effects of technological change, conceptually and empirically sound measures are crucial. In this chapter, I identified and discussed the shortcomings of existing measures and proposed new indices which address the problems and offer researchers flexible tools for analysing occupational task content.

The new indices entail the following improvements. First, both the routine-intensity and the complexity index have a clear theoretical interpretation: they capture the task characteristics that the RBTC and SBTC theories focus on, respectively. Second, the variables used to operationalise them really quantify the essence of the underlying concepts: repetitiveness and codifiability in the case of routine tasks, and higher-order skills such as effective communication, abstraction, and decision making in the case of complex tasks. This places some occupations very differently in the routine hierarchy, especially compared to the measures based on Autor & Dorn (2013). Third, the indices use survey data rather than expert-coded task data which is important for understanding within-occupation variation.

Furthermore, with the wave- and country-specific indices, a much more detailed analysis of the impact of RBTC and SBTC on occupational change becomes possible for the first time. This represents a significant improvement over the measures developed by Autor & Dorn (2013) and goes beyond the important contribution of Fernández-Macías & Hurley (2017). My descriptive analysis of the novel measures shows trends and differences that are too substantial to ignore. This calls into question the practice of applying task measures from one country and year, in very different contexts. Further research into the reasons for and consequences of these differences may have important implications for the understanding of occupational and technological change.

The main limitations of this approach to operationalising occupational task content are the sample size of the EWCS which prevents more disaggregated analyses at the occupation-sector level, and concerns about potential measurement error introduced by individuals understanding the survey questions differently. Furthermore, other authors have also proposed alternative approaches to Autor & Dorn (2013) and Fernández-Macías & Hurley (2017). For example, Salvatori (2018) develops a routine intensity index using UK data based on Autor and Dorn's (2013) methodology. However, Salvatori's and similar contributions lack a comprehensive methodological discussion and a generalisable and flexible set of measures. Thus, the measures proposed here will be useful for future research on the nature of contemporary technological and occupational change. Much remains to be learned about how work is changing and what role technology plays in the process.

In the context of this thesis, this chapter has answered the first overarching research question: how can occupational task content be measured in a way that corresponds to the underlying theoretical concepts and that accounts for the relevant dimensions of variation? To better align theory and measurement, I showed that indices based on a suitable set of questions from the EWCS yield a more realistic assessment of occupational routine intensity and can capture the idea that underpins SBTC with a measure of task complexity. The analyses have furthermore highlighted substantial variation within occupations, as well as over time and between countries that previous operationalisations have not accounted for. The analyses in the following chapters build upon the important insights from this chapter and use its measures to analyse which factors determine patterns of occupational change and wage hierarchies.

6

Routine-biased technological change does not always lead to polarisation: Evidence from 10 OECD countries, 1995 - 2013

Contents

6.1 Introduction

According to the routine-biased technological change hypothesis, in recent years, machines and software have taken over codifiable and repetitive tasks and made redundant many routine workers. Coupled with the argument (or often, as I shall argue, assertion) that these jobs are disproportionately found around the median of the wage distribution, this leads to a predicted pattern of job polarisation: a decline in relative employment in medium-wage routine jobs, accompanied by the growth of high-paying (cognitive non-routine) jobs and low-paying (manual non-routine) service occupations.[1] The widely accepted narrative is thus that RBTC is partly responsible for a growing polarisation of both employment and wages in rich countries. Yet, we are in danger of accepting an incomplete narrative of what happened during the last 30-or-so years.

While polarisation has occurred in some countries, many have seen upgrading instead, contradicting the original RBTC hypothesis. Thus, in this chapter, I use the task measures that I developed in chapter 5 and employment and wage data for 10 OECD countries between 1995 and 2013 from the Luxembourg Income Study to investigate the second research question of this thesis: can the diverse patterns of employment change in developed countries be reconciled with the key tenet of RBTC, that technological change everywhere substitutes for routine workers? I find that the answer is, yes, after another common assumption of the RBTC model is relaxed: not everywhere are routine occupations medium-wage. I therefore propose a refined theoretical model that explains how the relationship between routine tasks and occupational wages shapes whether employment change as a result of RBTC is upgrading or polarising. Here the routine-wage curves which I introduced in chapter 3.2 are employed to capture the relationship between routine tasks and occupational wages.

I argue that countries differ with regard to the shape of the routine-wage curve and that different patterns of employment change are primarily a reflection of this: occupational polarisation tends to occur only in countries with a hump-shaped RTI-wage curve where routine occupations cluster near the middle of the wage distribution. Where routine occupations command the lowest relative wages and the RTI-wage curve is monotonic, occupational upgrading follows. I go on to show that the hump shape is more prevalent in richer countries and that routine-wage schedules are not driven by non-standard work or compositional changes within occupations; moreover, they are relatively stable over time. This study thus bolsters the findings of a recent literature in economic sociology which calls into question a simplistic

[1] Throughout the chapter, polarisation is used in this way to refer to polarisation of employment in terms of occupational wages. Upgrading is used analogously. Unless otherwise noted, I refer to employment trends in terms of occupations rather than, say, social classes.

generalisation of the routine-bias hypothesis, but does so without abandoning the RBTC framework as such. Instead, it argues that the mechanism that has been correctly identified in the economics literature does not always lead to polarisation if country-specific characteristics are taken into account.

In the structure of the thesis, this chapter constitutes the heart of my theoretical argument. Here the key idea, that the same process of technological change may result in different patterns of employment change due to differences in relative occupational wages, is developed and substantiated empirically. Relying on the measures that I developed in chapter 5 as a crucial component of its empirical strategy, this chapter with its findings furthermore sets the stage for the third substantive research question in chapter 7, whether and how robotisation and LMI shape occupational wage hierarchies.

As I argued in chapter 2, both economists and sociologists have made crucial contributions to the literature on technological change and employment change. Economists have developed the RBTC framework in an attempt to better understand how ICT operates in the labour market. Autor et al. (2003) were the first to argue that computer capital substitutes for workers in performing routine tasks and complements them in nonroutine tasks. Later research found routine occupations to be concentrated around or just below the median of the wage distribution in the UK and in the US (Autor et al. 2008, Goos & Manning 2007). Based on these observations, RBTC was taken to predict widespread polarisation of the employment structure in the US and beyond, which has been corroborated by empirical studies such as Goos et al. (2014) and Michaels et al. (2014).

Sociologists generally acknowledge the importance of technology and the presence of employment polarisation, at least in the US (Wright & Dwyer 2003). Even where they do not dispute polarisation, however, many sociologists criticize the disregard in the RBTC literature for other macro trends such as the emergence of the "care economy" or immigration (Dwyer 2013, Oesch 2015). Yet, these developments affect the demand for and supply of labour largely beyond the direct effect of technology. Likewise, labour market institutions receive little attention from scholars of RBTC despite their prominence in the inequality literature (Koeniger et al. 2007, Mishel, Shierholz & Schmitt 2013, Parolin 2021). Others again maintain that skill-bias is still a more accurate account of technological change (Oesch & Piccitto 2019). In light of these factors, pervasive polarisation across countries appears unlikely.

A multitude of empirical studies thus investigate this issue. This literature can be summarised as follows. Economists generally agree that there has been widespread polarisation of the employment structure, although there are different views to which degree this is attributable to RBTC. Sociologists, on the other hand, contest the pervasiveness of polarisation, and where they do find evidence for it, question whether RBTC is the only – or even the primary – explanation. The more nuanced sociological theories therefore present a serious challenge to simplistic accounts of technological change. With few notable exceptions (e.g. Fernández-Macías & Hurley 2017), the technical assumptions behind the RBTC model have not been the focus of this literature.

This chapter, on the other hand, argues that even in the absence of an effect from the aforementioned factors, polarisation would not always follow RBTC. Its argument is that technological change is indeed biased against routine occupations, but that the rigid assumption that routine occupations are medium-wage has prevented the RBTC model from providing a generalisable account of recent employment trends.

6.2 A refined model of routine-biased technological change

The refined model that I propose in this chapter takes issue with two central claims of the RBTC literature. First, I challenge the notion that RBTC has supplanted SBTC. As sociologists have previously recognised, SBTC and RBTC are not alternatives; rather they are related yet distinct processes that may unfold simultaneously and with varying intensity in different countries (Oesch 2013, Oesch & Rodriguez Menes 2011). My model incorporates this insight. Secondly, the assumption that countries' routine-wage schedules are effectively identical does not hold: routine occupations are not everywhere medium-wage (Fernández-Macías & Hurley 2017), but neither are they always near the bottom, and thus the same technological change may lead to different outcomes based on a country's existing wage hierarchy. Incorporating these insights yields a more encompassing and realistic account of technological and employment change. As I show in the empirical section below, it can explain most of the deviations from the standard RBTC model, but unlike some of the criticisms in the sociological literature, it does so while retaining the basic logic of the framework.

RBTC and SBTC are not mutually exclusive processes: they operate on different occupational characteristics and may have distinct effects on the employment structure. Yet, economists often argue that RBTC has supplanted SBTC, while some sociologists consider lack of polarisation evidence in favour of SBTC over RBTC. This either-or approach does not seem warranted. Instead, I argue that skill-biased upgrading is likely to continue alongside RBTC. For example, Caines et al. (2017) show that in the US technological change between 1980 – 2005, and hence employment and wage growth, was biased towards complex tasks in an apparent corroboration of the SBTC hypothesis beyond its original core period in the 1980s. Balsmeier & Woerter (2019) report similar findings for Swiss firms in recent years.

Furthermore, many more recent additions to the OECD have not yet reached the level of economic maturity of, say, the US or Germany. Hence, economies like Chile and the Eastern European countries still had some catching up to do in the 1990s and 2000s, which constitute the core period of my analyses. A broader theory of technological change and employment must be able to accommodate these countries – developed but not at the technological frontier – as well. This study is one of the few to include these countries, and insofar as SBTC usually precedes RBTC, we should expect them to experience SBTC during the period of analysis. These considerations suggest that SBTC should not be discarded as part of a general theory of technology and labour market dynamics. I thus follow Oesch (2013) who is one of the few scholars who explicitly investigate SBTC alongside RBTC, but with a larger and more diverse group of countries.

The argument that RBTC results in occupational polarisation is based on the empirical observation that in the US and UK around 1980, routine occupations tended to be medium-wage (Autor et al. 2008, Goos & Manning 2007). Scholars such as Goos et al. (2014) have extended this assumption to large and diverse samples of countries. However, there is no iron law which dictates that medium-wage occupations are particularly routine-intensive; unlike with SBTC, there is no strong theoretical expectation where routine occupations should be in the wage hierarchy (Oesch & Piccitto 2019). Indeed, using refined and updated measures of routine-intensity Fernández-Macías & Hurley (2017) have found that on average in Europe, low-paying occupations tend to be the most routine-intensive, thus directly contradicting Goos et al. (2014).

This indicates that the hump-shaped routine-wage curve cannot be as ubiquitous as assumed in the labour economics literature. However, Fernández-Macías & Hurley (2017) do not look at routine-wage distributions in individual countries and it remains an open empirical question how similar they are. That they are identical appears unlikely: Eurofound (2014) shows that there is a significant degree of variation in occupational wage rankings in Europe. It is therefore possible that RBTC primarily affects low-wage occupations in one country and medium-wage occupations in another, resulting in different employment trends when occupations are ranked by their average wage. Hence, a generalisable theory of RBTC should not rely on a common routine-wage schedule across countries for its predictions. Yet, to my knowledge, no other study explicitly takes this into account.

My refined theory of RBTC can be visualised as in figure 6.1. It assumes that an exogenous force, for example a decline in the price of computer capital, leads to both a rise in the demand for workers performing complex tasks (SBTC) and lower demand for routine workers (RBTC). The key word here is "and", as previous theories usually focus on only one of these avenues. The crucial novel element is the explicit consideration of the task-wage relationship on a country-by-country basis. In all countries, complexity increases monotonically with the average occupational wage, as depicted on the left of panel B. However, the same is not true for RTI and wages. In some countries, RTI is strictly decreasing with increasing wages as posited by Fernández-Macías & Hurley (2017, middle of panel B), while in others,

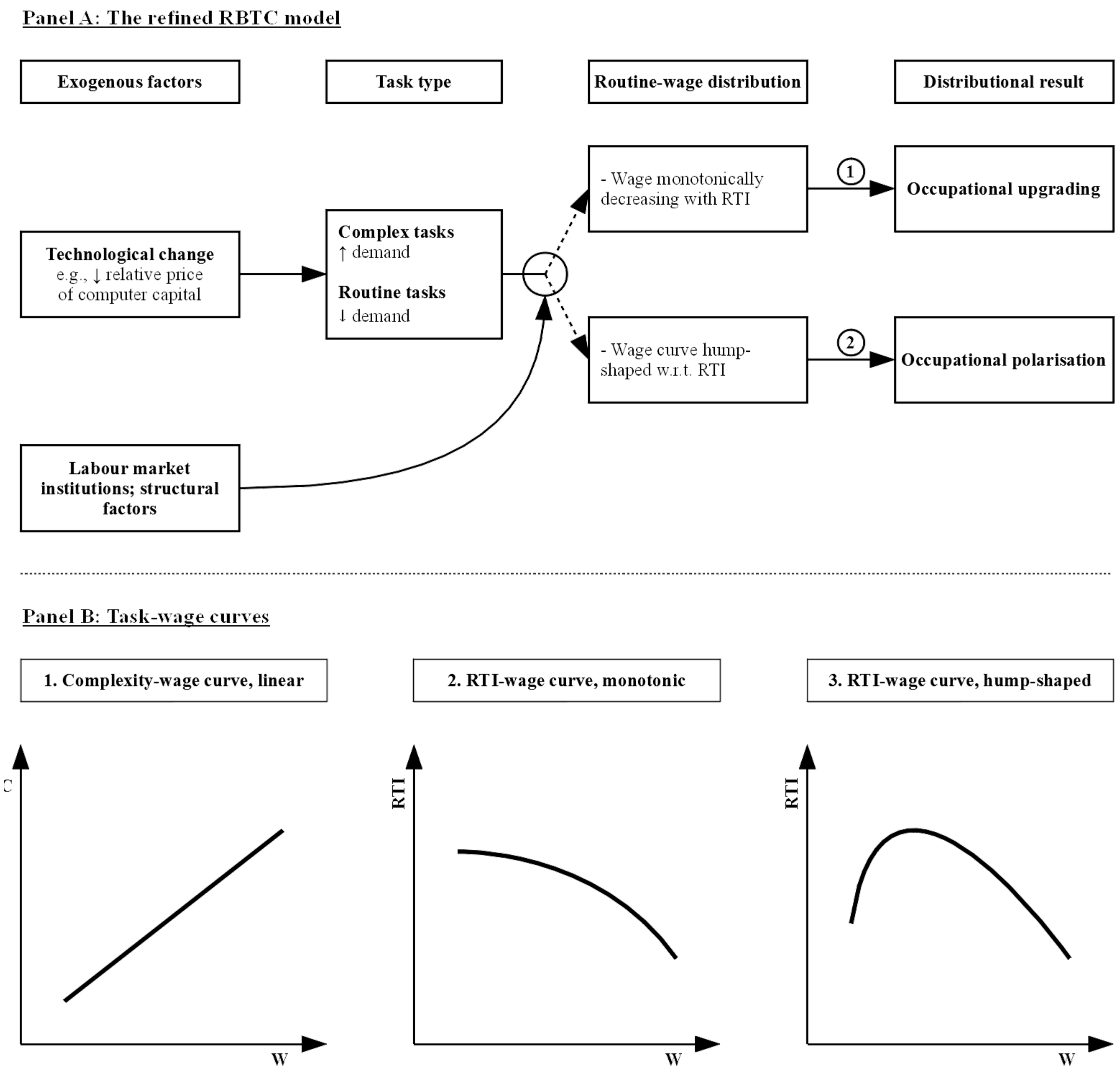

Figure 6.1: *Panel A provides a simple visualisation of the refined RBTC model. Panel B shows the stylised task-wage curves. Panel B1 shows the complexity wage schedule, panel B2 the routine wage schedule in scenario 1 of panel A, and panel B3 shows the routine wage schedule of scenario 2 in panel A under which polarisation occurs.*

the curve is hump-shaped with the highest RTI values for medium-wage occupations as in Autor et al. (2003, right of panel B).

Whether a country experiences upgrading or polarisation thus crucially depends on the routine-wage curve: in scenario 1 of panel A, the complexity-wage curve and the monotonic RTI-wage schedule combine to result in occupational upgrading. Due to the negative correlation between complexity and RTI, SBTC and RBTC reinforce each other, and we see a reallocation from low-wage, low-complexity, high-routine employment to high-paying, complex, non-routine jobs. This scenario is in line with the findings of Fernández-Macías & Hurley (2017). Conversely, polarisation occurs in scenario 2, when the complexity-wage curve is combined with the hump-shaped RTI-wage curve. This is similar to the standard assumption in labour economics (e.g. Goos et al. 2014), although here the continuing importance of skill-bias in technological change is acknowledged. This approach recognises the value of both the routine-bias and skill-bias models and enriches them with insights from the sociological literature. It thus provides a more realistic account of technological and employment dynamics.

Of course, technology is not the only factor that influences employment trends, as the studies discussed above show (see, e.g., Fernández-Macías & Hurley 2017, Oesch 2015, Wright & Dwyer 2003). Structural factors such as female and high-skilled labour supply and institutions such as labour unions may have a direct impact on the occupational structure. In the empirical investigation, country-occupation and study-wave fixed-effects will pick up much of the effect of these factors. Importantly for my theory, structural and institutional differences may also play an important role in determining wage hierarchies and may thus be the underlying reason behind the different routine-wage curves. The sociological critiques of the RBTC theory are therefore accounted for in this refined model. An in-depth empirical investigation of these links is the subject of the next chapter of this thesis.

6.3 Data and descriptive analyses

6.3.1 Employment and labour income

The main analytical sample for this study covers 10 countries during the period 1995 – 2013: Chile, Czech Republic, Finland, Germany, Hungary, Luxembourg, Netherlands, Slovenia, Spain, and the United States. I use employment and wage data from the Luxembourg Income Study.[2] I calculate occupational employment shares and average wages at the ISCO-88 2-digit level for 167 country-years covering 27 countries and spanning the time period from 1974 until 2016. From this larger dataset, the sample is selected to include the 10 countries with an uninterrupted

[2] For more details on the LIS, see chapter 4.

time-series for LIS waves IV to IX. The main sample includes all workers, with supplementary analyses looking at working-age full-year full-time workers only.

I prepared this dataset by calculating occupational employment shares and average annual wages at the most detailed level of the national occupational classification recorded in the LIS database. Subsequently, I used the crosswalks provided by Mahutga et al. (2018) to harmonise the data at the 2-digit ISCO-88 level. In doing so, I added 28 additional country-years that were not included in the dataset of Mahutga et al. (2018) and corrected some minor coding mistakes. To my knowledge, this constitutes the first comparative dataset of occupational employment shares and annual wages at the 2-digit ISCO-88 level based on LIS data. As discussed in chapter 4, the use of annual labour income has the drawback that I cannot account for differences in hours worked, but where possible I use hourly wage data to ensure that my findings are not driven by this.

6.3.2 Task content data

I use the indices of routine task intensity and task complexity that I developed in the previous chapter (Haslberger 2021*a*). As I explained there, these measures have several advantages over the widely used RTI index from Autor & Dorn (2013), and offer further improvements over the approach developed in Eurofound (2014) and Fernández-Macías & Hurley (2017). Routine tasks are defined as tasks that are repetitive and codifiable (Fernández-Macías & Hurley 2017), while the complexity index is based on Caines et al. (2017) who define complex tasks as those requiring higher-order skills such as effective communication, abstraction and decision making.

I use an index based on data from the EU-15 countries throughout the study; however, using country-specific measures does not substantively change the conclusions for European countries. Applying measures based on an entirely European sample to the US and Chile is certainly not unproblematic. However, compared to the prevailing practice of using an index based on the American Dictionary of Occupational Titles from 1977 for all countries and periods, my approach undoubtedly promises more robust insights. Indeed, I showed in chapter 5 that in most cases, a pooled index will rank occupations in a country more similarly to that country's actual ranking than applying the ranking of another individual country in the sample. It therefore appears highly likely that such a pooled measure will also exhibit better out-of-sample performance than a measure from any one country-year. Since the task data are measured between 2000 and 2015, and thus do not fully overlap with the period of analysis, I use task values that are averaged over the entire period, rather than wave-specific measures.

6.3.3 Employment trends by occupational wage

This section presents data on employment trends by occupational wage and establishes that there is indeed an empirical puzzle: there is no universal pattern of employment change that conforms to a simple reading of either the SBTC or RBTC theory. Panel A of table 6.1 provides an overview similar to that in Goos et al. (2014), with occupations ordered and divided roughly into terciles based on their average occupational wage. Unlike Goos et al. (2014), I include all occupations whereas they exclude agricultural occupations and what appear to be occupations with heavy public sector involvement (ISCO groups 11, 23, 33, 61, and 92).[3] Comparing my table to that in Goos et al. (2014), we find that the ranking of occupations in terms of their average wage is highly similar: the rank-order correlation with the 21 2-digit occupations that are included in Goos et al. (2014) is 0.98.

While the ordering in terms of wages is almost identical, there are substantial differences in routine scores and employment trends. Above all, middling occupations do not stand out as disproportionately high-routine in my data. It appears that the statement, "routine occupations cluster around the middle of the wage distribution" is, at least at this level of generality, misleading.[4] My data furthermore show

[3]The excluded occupations cumulatively account for almost 10 percent of total employment. Senior officials and teaching professionals fall into the highest-earning tercile; the other three groups are among the lowest-paid. While there may be valid reasons for this exclusion, they are not made explicit in Goos et al. (2014).

[4]In Goos et al. (2014), the discrepancy between middling and low-paying occupations in terms of RTI is inflated by classifying the high-routine groups 82 and 74 as middling occupations, despite them belonging to the lowest-paying third. High-routine group 81 is classified as a middling occupation in Goos et al. even though it is in the top-earning tercile. This swells the medium-wage group to comprise the 22nd to 68th percentiles of occupational wages. It is only by using cut-off

that what mainly distinguishes middling from low-wage occupations is their higher average complexity, while they are fairly similar in terms of routine-intensity. It therefore seems likely that greater complexity explains much of the wage premium in middling occupations. Most importantly, table 6.1 provides no evidence for pervasive employment polarisation if we use the full set of occupations. Middling and low-paying occupations both saw sizeable employment reductions, while high-paying occupations were the only group to expand their employment share in 2013 compared to 1995. This finding is robust to looking at quintiles or more arbitrarily defined groups as in Goos et al. (2014); in fact, with quintiles the average changes look even more clearly upgrading.[5] However, further analyses, reported in table C.1 in appendix C.2, reveal a gendered pattern of (tentative) upgrading for women and polarisation for men, as recorded in Eurofound (2014).[6] On average, these results therefore suggest that occupational change was largely upgrading in terms of wages.

If we instead plot the trajectory of high-, medium-, and low-wage occupations in individual countries, we see a diversity of patterns that is consistent with my refined theory. In figure 6.2 below, I ranked occupations by their average wage in the respective country in the first period for which I have data and then classified them so as to represent as close as possible to one third of employment.[7] The graphs indicate that neither the upgrading nor the polarisation narrative can explain

points in a strategic manner and excluding the aforementioned occupations that Goos et al. achieve the stark contrasts in their analysis. Moreover, the high routine-intensity of middling occupations in Goos et al. is driven largely by the implausibly high RTI score for office clerks who, as previous research has argued, have seen their range of tasks expand to be a far cry from the high-routine occupation of the past (see also chapter 5).

[5] See figure C.1 in appendix C.1.

[6] Table C.1 in appendix C.2 shows a simple model of employment change regressed on log annual wages and the share of full-year, full-time (FYFT) workers in 1995 which suggests a somewhat ambiguous picture. Just looking at the wage variables, employment changes appear more polarising than upgrading. Running the models separately for women and men, we see that the overall polarisation is driven by men while for women there is a statistically insignificant tendency towards upgrading. If we control for the FYFT share, the quadratic relationship gives way to a linear one, but since data on status in employment are only available for 5 of the 10 countries, the results are not directly comparable. An important takeaway from these regressions is that the coefficient on the share of FYFT workers is negative and statistically significant (except for women only), which suggests that non-standard employment has expanded more strongly than standard employment. The results for women and men separately are similar to the 10-country model. In my data, therefore, we see a gendered pattern of (tentative) upgrading for women and polarisation for men, as recorded in Eurofound (2014).

[7] With 26 ISCO codes, tercile shares even in the initial period are never exactly 33.3 percent.

employment trends across the developed world. The United States and Germany are textbook examples of employment polarisation, as defined in chapter 1.2.1. Employment in high-wage occupations increased strongly and low-wage employment grew moderately in both countries, at the expense of medium-wage occupations. Yet, only Luxembourg and The Netherlands also experienced growing high-wage and low-wage employment at the expense of middling occupations.

In five of the remaining six countries, both medium- and low-wage employment shares declined, resulting in an upgrading employment structure. Chile and Spain experienced relatively straightforward upgrading of the occupational structure. These countries, of course, are latecomers to the club of liberal democracies and as such had to catch up economically to countries like the US. Czech Republic and Slovenia also experienced a strongly upgrading pattern since the 1990s, the recent dip in high-wage employment largely being an artefact of the change to ISCO-08. Again, this appears consistent with a story of economic catch-up after democratisation. In Finland, low-wage employment was marginally lower in 2013 than in 1995, while medium-wage employment declined more substantially and high-paying jobs expanded, in a pattern consistent with polarised upgrading. Hungary is the only outlier in that the data show occupational downgrading. This is in line with Eurofound (2017) who also find downgrading in Hungary from 2011 – 2016; my findings suggest that this was in fact a more durable process.

Table 6.1: Average wage, routine, and complexity rankings.

PANEL A: WAGE RANKING						PANEL B: RTI RANKING					PANEL C: COMPLEXITY				
ISCO	E 1995	E 2013	Δ%	RTI	COMP.	ISCO	RTI	E 1995	E 2013	Δ%	ISCO	COMP.	E 1995	E 2013	Δ%
12	3.19	3.01	-0.18	-0.22	0.47	11	-0.34	0.41	0.43		21	0.61	3.30	4.88	
21	3.30	4.88	1.58	-0.19	0.61	23	-0.33	4.23	4.23		24	0.49	3.82	5.71	
11	0.41	0.43	0.02	-0.34	0.49	33	-0.32	1.15	1.55		11	0.49	0.41	0.43	
22	1.66	2.32	0.66	-0.07	0.38	24	-0.23	3.82	5.71		31	0.48	3.37	3.02	
24	3.82	5.71	1.88	-0.23	0.49	12	-0.22	3.19	3.01		12	0.47	3.19	3.01	
23	4.23	4.23	-0.01	-0.33	0.23	21	-0.19	3.30	4.88		22	0.38	1.66	2.32	
31	3.37	3.02	-0.35	-0.07	0.48	34	-0.17	7.43	8.77		34	0.34	7.43	8.77	
13	3.86	3.34	-0.53	-0.09	0.07	13	-0.09	3.86	3.34		41	0.26	9.45	9.06	
34	7.43	8.77	1.32	-0.17	0.34	31	-0.07	3.37	3.02		**High**	**0.39**	**32.64**	**37.2**	**4.57**
81	1.31	1.15	-0.16	0.30	-0.20	22	-0.07	1.66	2.32		23	0.23	4.23	4.23	
High	**32.58**	**36.87**	**4.29**	**-0.16**	**0.35**	**Low**	**-0.19**	**32.43**	**37.26**	**4.84**	32	0.19	2.82	3.61	
72	6.25	5.25	-1.01	0.13	0.04	52	-0.05	5.24	5.26		33	0.15	1.15	1.55	
83	4.51	4.23	-0.28	0.06	-0.41	41	-0.02	9.45	9.06		13	0.07	3.86	3.34	
73	1.24	0.90	-0.34	0.23	-0.08	32	-0.02	2.82	3.61		72	0.04	6.25	5.25	
32	2.82	3.61	0.79	-0.02	0.19	51	0.03	7.17	9.17		42	0.01	2.61	2.77	
41	9.45	9.06	-0.42	-0.02	0.26	83	0.06	4.51	4.23		73	-0.08	1.24	0.90	
71	5.57	4.49	-1.09	0.27	-0.12	61	0.08	2.01	1.56		71	-0.12	5.57	4.49	
33	1.15	1.55	0.39	0.17	0.01	**Med.**	**0.00**	**31.2**	**32.89**	**1.69**	81	-0.20	1.31	1.15	
42	2.61	2.77	0.15	0.39	-0.29	72	0.13	6.25	5.25		51	-0.22	7.17	9.17	
Med.	**33.61**	**31.86**	**-1.75**	**0.08**	**0.02**	91	0.14	6.61	6.80		**Med.**	**-0.01**	**36.21**	**36.46**	**0.26**
82	3.92	3.45	-0.47	-0.32	0.15	42	0.17	2.61	2.77		74	-0.28	3.14	1.79	
74	3.14	1.79	-1.35	0.34	-0.28	92	0.21	2.23	0.62		82	-0.29	3.92	3.45	
51	7.17	9.17	1.99	0.03	-0.22	73	0.23	1.24	0.90		52	-0.30	5.24	5.26	
61	2.01	1.56	-0.46	0.08	-0.35	93	0.25	3.49	2.61		61	-0.35	2.01	1.56	

Table 6.1 continued from previous page

	PANEL A: WAGE RANKING					PANEL B: RTI RANKING					PANEL C: COMPLEXITY				
ISCO	E 1995	E 2013	Δ%	RTI	COMP.	ISCO	RTI	E 1995	E 2013	Δ%	ISCO	COMP.	E 1995	E 2013	Δ%
93	3.49	2.61	-0.88	0.25	-0.49	71	0.27	5.57	4.49		83	-0.41	4.51	4.23	
52	5.24	5.26	0.01	-0.05	-0.30	81	0.30	1.31	1.15		93	-0.49	3.49	2.61	
91	6.61	6.80	0.18	0.14	-0.63	74	0.34	3.14	1.79		92	-0.56	2.23	0.62	
92	2.23	0.62	-1.62	0.21	-0.56	82	0.39	3.92	3.45		91	-0.63	6.61	6.80	
Low	**33.81**	**31.27**	**-2.54**	**0.14**	**-0.38**	**High**	**0.23**	**36.37**	**29.85**	**-6.52**	**Low**	**-0.43**	**31.16**	**26.33**	**-4.82**

Note: "ISCO" refers to the 2-digit ISCO-88 occupational codes. "E 1995" and "E 2013" denote the employment share in percent in the respective year. Δ% is the percentage change in employment shares. The wage ranking is based on pooled average wages in 1995.

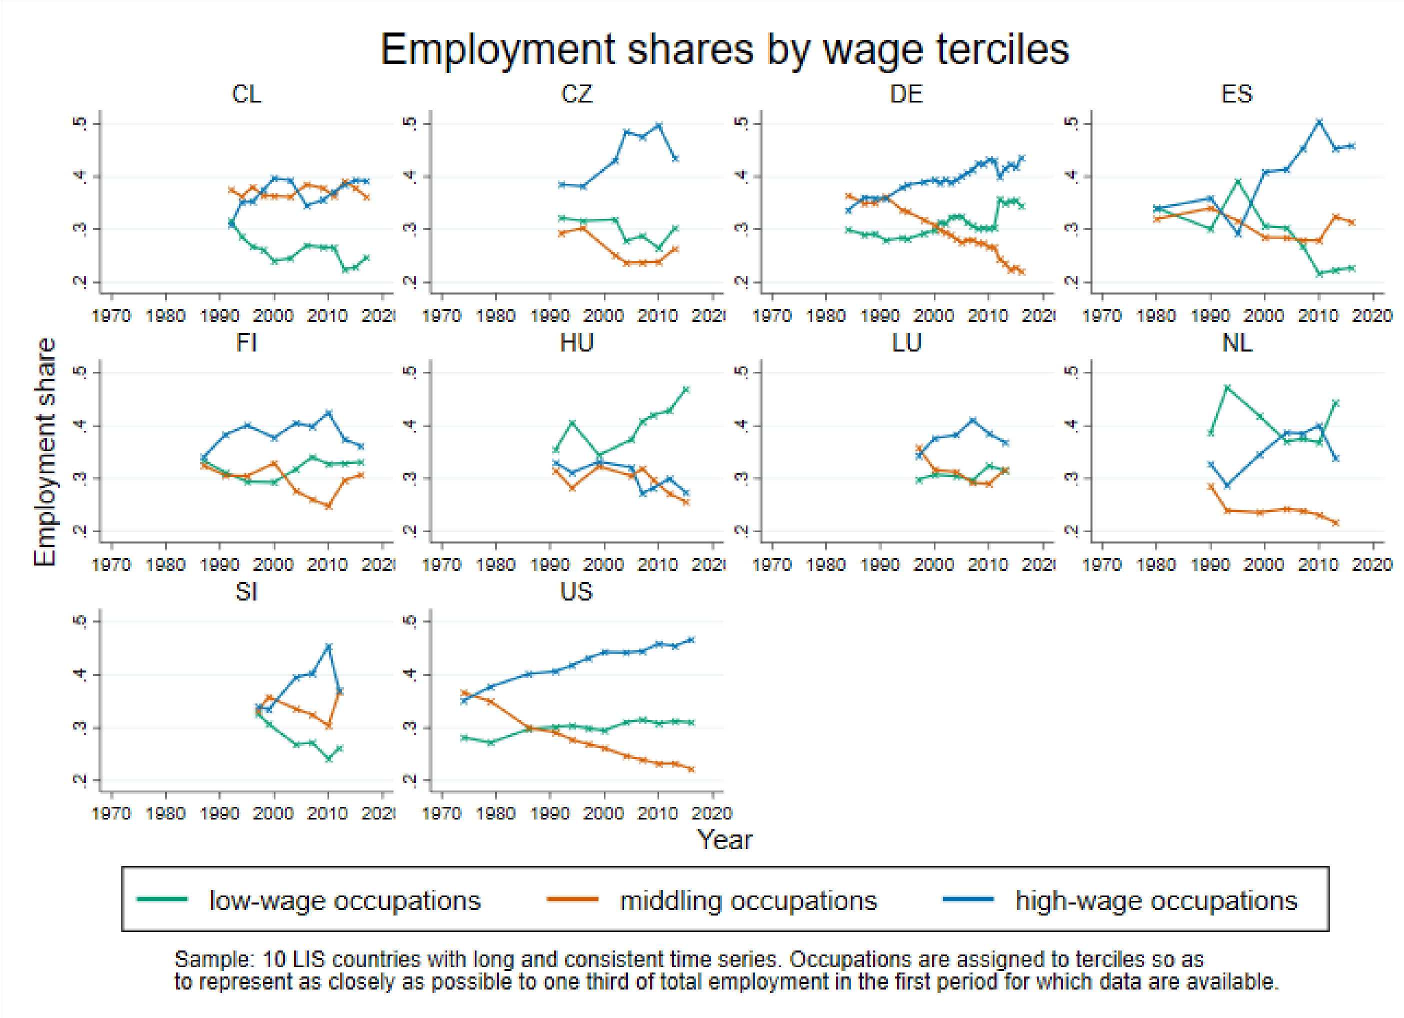

Figure 6.2: *Employment changes by wage tercile in 10 LIS countries.*

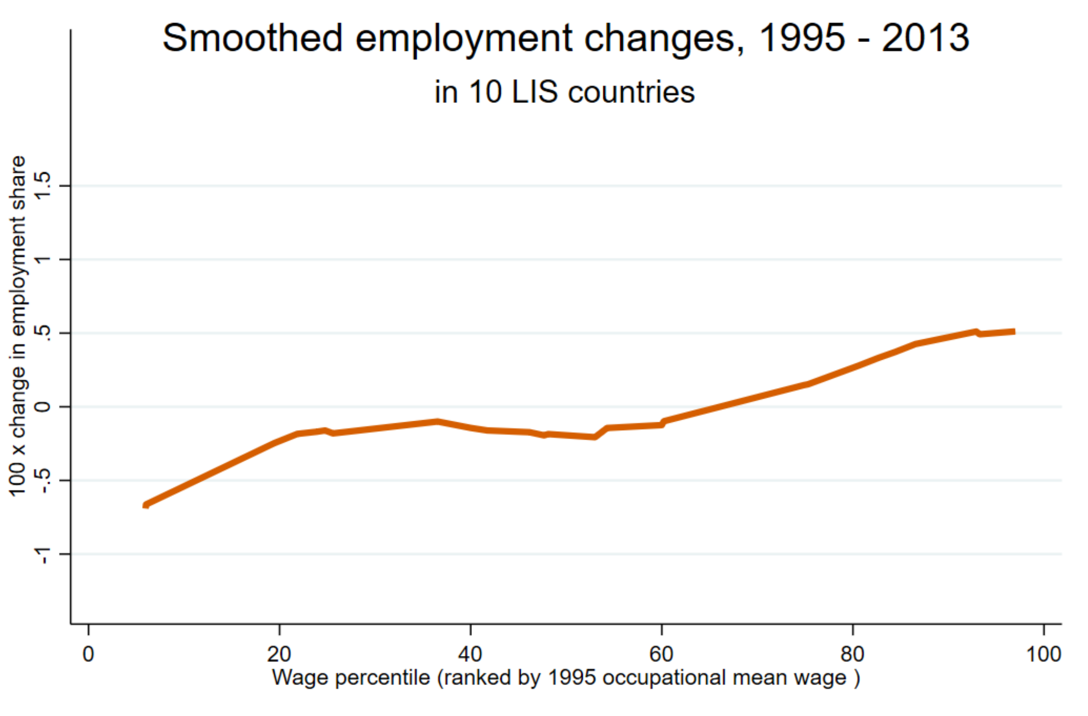

Figure 6.3: *Smoothed changes in employment shares based on a locally weighted regression with bandwidth = 0.75.*

The overall picture is compounded by the change from ISCO-88 to ISCO-08 that happened in many European countries around 2010. Despite the use of the official crosswalk in Mahutga et al. (2018), a significant discontinuity cannot be avoided and we see a sudden dip in the employment share of high-wage occupations accompanied by a jump in the share of low-wage occupations. The example of Germany, where we have annual data, shows that this is not a trend but a one-off effect, as previous trends continue afterwards. The findings for Hungary, and to a lesser extent Spain, Luxembourg, and the Netherlands, may furthermore be more volatile due to low sample sizes in some waves. For example, in Spain the 1995 wave includes only 1,144 respondents, compared to over 10,000 in later waves, and shows a marked jump compared to surrounding waves. The overall trend of occupational upgrading, however, is unaffected by this.

Figure 6.3 shows the average pattern of employment change in my sample, based on a locally weighted regression. The figure provides a neat illustration of the contemporaneousness of upgrading and polarisation. On average, relative employment expanded in roughly the best-paid third of occupations. The greatest declines in relative employment were recorded in the bottom quintile. This points towards upgrading, however, embedded in the overall upgrading trend is a small dip around the 50$^{\text{th}}$ wage percentile. Figure C.2 in appendix C.1 shows that this average curve is the result of diverse country trends, in line with the findings from figure 6.2.

To summarise, the analyses presented here suggest that occupational upgrading was the most common pattern in a range of developed countries, with a sizeable minority of countries experiencing polarisation – among them the two largest and best-studied economies in the sample, Germany and the US.[8] RBTC as commonly understood cannot explain these diverse experiences, so there is indeed a puzzle that my refined theory can address.

6.3.4 Continuing SBTC in the data

This section presents evidence for the first element of my refined model: simultaneous skill-biased and routine-biased technological change. Independent of where in the wage distribution they are located, RBTC requires that routine occupations decline and low-routine occupations grow, and SBTC entails that complex occupations expand while simple occupations contract. In panels B and C of table 6.1, we see that employment change between 1995 and 2013 was indeed highly skill- and routine-upgrading. Occupations are ordered by their task requirements, from least to most routine-intensive in the panel in the middle, and from most to least complex in the panel on the right. Occupations scoring low on RTI or high on complexity increased their cumulative employment share by almost 4.8 and 4.6 percent, respectively.

Total employment in occupations with average RTI and complexity scores did not change much, with small increases of 1.7 and 0.3 percent, while the most routine-intensive and the least complex occupations lost almost 6.5 and 4.8 percent of cumulative employment, respectively. Of course, many simple occupations are

[8] If the appraisal of employment changes is limited to the period 1995 – 2013 and occupations are allocated to terciles based on their wages in 1995, Finland exhibits a polarising pattern as well while the trend in the US can be described as weak polarisation, with strong employment growth in high-wage occupations, strong declines in middling occupations, and essentially constant employment in low-wage occupations (see figure C.3 in appendix C.1).

also routine-intensive, but there are several exceptions to this pattern which show that the measures are only partly collinear. The linear patterns largely hold if we divide occupations more finely into quintiles, as can be seen in appendix C.1. This finding is similar to Oesch (2013) who also finds evidence for both RBTC and SBTC if occupations are divided into task and skill terciles. However, he concludes that employment change has corresponded to RBTC in Britain and Spain and to SBTC in Denmark, Germany, and Switzerland, while this study argues that RBTC, properly understood, can account for employment changes in upgrading countries as well.

Simple regression analyses provide further support for a linear relationship between tasks and employment shares. In table 6.2, I regress occupational employment shares on the RTI and complexity measures, first for all workers and FYFT workers and then for women and men separately. I interact the task measures with a linear time-trend to deal with the fact that they are constant over time. All regressions also include country-occupation and year fixed-effects.[9] With this, I follow Goos et al. (2009, 2014) who use a similar approach to estimate the effect of task content on labour demand. My results in panel A largely confirm the findings of Goos et al. (2014). Routine-intensity has a stable and highly statistically significant negative association with occupational employment (column 1). The coefficient on task complexity is positive and marginally statistically significant if entered alone (column 2); however, in a "horse-race", the coefficient on RTI retains its size and statistical significance, while task complexity loses both (column 3). Running the models for FYFT workers only leads to very similar results, albeit less precisely estimated. It appears, therefore, that the trends for non-standard workers were not fundamentally different than for FYFT workers, alleviating some of the concerns articulated above.[10] Running the models separately for women and

[9]Besides institutional and supply-side factors, the fixed-effects may pick up some of the effect of technology, making the coefficients on RTI and complexity lower-bound estimates. I owe this point to an anonymous reviewer.

[10]Table C.2 in appendix C.2 extends the analyses in table 6.2 by expanding the sample in panels A and B to include all available LIS country-years between 1995 and 2016. The results are in line with, and in fact slightly stronger than, in the sample presented here. In panel C, I control for the share of FYFT workers in each occupation in the countries for which this information is available. This analysis shows significant results where table 6.2 does not. Again, therefore, the relationship between RTI and occupational employment shares is robust and appears not to be

Table 6.2: Tasks and employment demand

PANEL A: BY WORKER STATUS						
	ALL WORKERS			FYFT WORKERS ONLY		
DV: Employment share	(A1)	(A2)	(A3)	(A4)	(A5)	(A6)
Time-trend interacted with:						
RTI	-0.84***		-1.08***	-0.62**		-0.6
	-0.31		-0.41	-0.31		-0.37
Complexity		0.21*	-0.19		0.22	0.02
		-0.11	-0.13		-0.16	-0.19
Observations	1,558	1,558	1,558	1,558	1,558	1,558
Country-occupations	260	260	260	260	260	260
R-squared	0.077	0.029	0.083	0.076	0.068	0.076
PANEL B: BY SEX						
	WOMEN			MEN		
DV: Employment share	(B1)	(B2)	(B3)	(B4)	(B5)	(B6)
Time-trend interacted with:						
RTI	-1.19**		-1.56**	-0.61**		-0.67*
	-0.48		-0.67	-0.28		-0.38
Complexity		0.25	-0.28		0.22*	-0.05
		-0.19	-0.27		-0.12	-0.15
Observations	1,533	1,533	1,533	1,555	1,555	1,555
Country-occupations	260	260	260	260	260	260
R-squared	0.073	0.035	0.079	0.059	0.044	0.059

Note: All point estimates (and standard errors in parentheses) are multiplied by 100. Robust standard errors are clustered by country-occupation. All occupations are weighted by their initial employment share in 1995. All regressions include country-occupation fixed-effects and wave fixed-effects. Sample: 10 LIS countries 1995 – 2013. *** p<0.01, ** p<0.05, * p<0.1

men in panel B yields similar findings; interestingly, the estimated coefficient sizes are larger and more robust for women. Hence, we can conclude that RTI and task complexity are associated with employment trends in the ways my theory would predict. The regression results suggest, however, that the effect of routine-intensity is the quantitatively more important one.

6.4 The task-wage schedules

The crucial novelty in my refined theory is the mechanism that leads from the observed patterns of technological change to the observed patterns of employment

driven by differences in non-standard work.

change in terms of wages. We have just seen that technological change has been both skill-biased and routine-biased in recent years. This section first confirms that the relationship between occupational complexity and wages is indeed monotonic and positive in all countries. It then shows that there are two distinct types of relationship between occupational wages and RTI: one linear or monotonic, as found in Fernández-Macías & Hurley (2017), and the hump-shaped relationship that has formed the backbone of the traditional RBTC theory as in Goos et al. (2014), as depicted in panel B of figure 6.1 above. I also show that the shape of the task-wage curves within countries is relatively constant over time, even as individual occupations sometimes change their position.

Since the different shapes of the distributions are best illustrated graphically, I fit fractional-polynomial prediction plots for each of the 163 country-years which provide employment and wage data in my LIS dataset. I use the *fpfitci* command in Stata in its default configuration, that is, a second-degree polynomial estimated with *regress*. Occupations are weighted by their employment share in the respective country-year. The use of a fractional polynomial function serves to better detect nonlinearities where they arise, while remaining relatively agnostic regarding the functional form of the relationship. We can thus determine whether a relationship is linear or not without having to fit different equations and comparing fit statistics. I show that both patterns do indeed arise, and that the two stylised patterns in panel A of figure 6.1 capture the overwhelming majority of cases in my dataset.[11] Note that the task measures are calculated from data for the period 2000 – 2015, while LIS employment data extend further back in many countries. Yet, this is an assumption that has to be made if we want to move away from using DOT data that predates the LIS data and is based on the US only – surely a much stronger assumption than the one made here.

[11] Deviations from the proposed patterns appear to be due mostly to measurement error in very small occupations where wages may be measured imprecisely.

6.4.1 The complexity-wage schedule

We start with the complexity-wage schedule, which is the more straightforward of the two task-wage schedules. Human capital theory predicts that wages increase linearly, by and large, with complexity or skill demands (see, e.g., Murphy & Topel 2016). There is indeed a large literature which shows that workers who perform complex tasks are rewarded in the labour market, and that more complex jobs require a higher level of skills, usually proxied by education (see, e.g., Goldin & Katz 2008). There may be some exceptions to this pattern related to gender issues; for example, there is some evidence that jobs which require medium levels of skills and training but are traditionally associated with women tend to have lower average wages than lower-skilled jobs which are traditionally associated with men (Tahlin 2019). Overall, however, that more complex jobs command higher wages, is uncontentious.

Figure 6.4 investigates this line of argument and illustrates the complexity-wage curve in my 10-country sample in LIS wave V, which corresponds to the year 2000. For reasons of space, I cannot present the complexity-wage schedules for all available 163 country-years here, but detailed results are available upon request. My analyses show precisely what theory predicts: there is a clear positive relationship between complexity and wages. Moreover, this relationship is approximately linear in most countries, and it is monotonically increasing in all countries except the US. There, it is only the exceptionally low pay of teaching associate professionals that introduces a nonmonotonicity at the bottom of the distribution; if ISCO group 33 is excluded the complexity-wage curve is monotonically increasing in the US as well.[12]

The remaining panels of figure 6.4 show that the postulated relationship between task complexity and wages holds in a wide range of circumstances. For example, in

[12] I suspect that this has to do with the official crosswalks from the US SOC classifications to ISCO. The only SOC groups partly assigned to ISCO group 33 are childcare workers (4600), teacher assistants (2540), and other education, training and library workers (2550) – in most other classifications at least some actual teachers are classified as associate teaching professionals, boosting average wages of the group. Interestingly, the position of group 33 in the US wage hierarchy does not change if we look at full-year full-time workers and hourly wages only. In a few other cases the linear relationship is distorted by an occupation which is assigned an implausibly low or high average wage – in some cases more than an entire log point from the next lowest- or highest-earning occupation. Such differences are arguably due to measurement error in small samples and should not be seen to undermine the general pattern in the data.

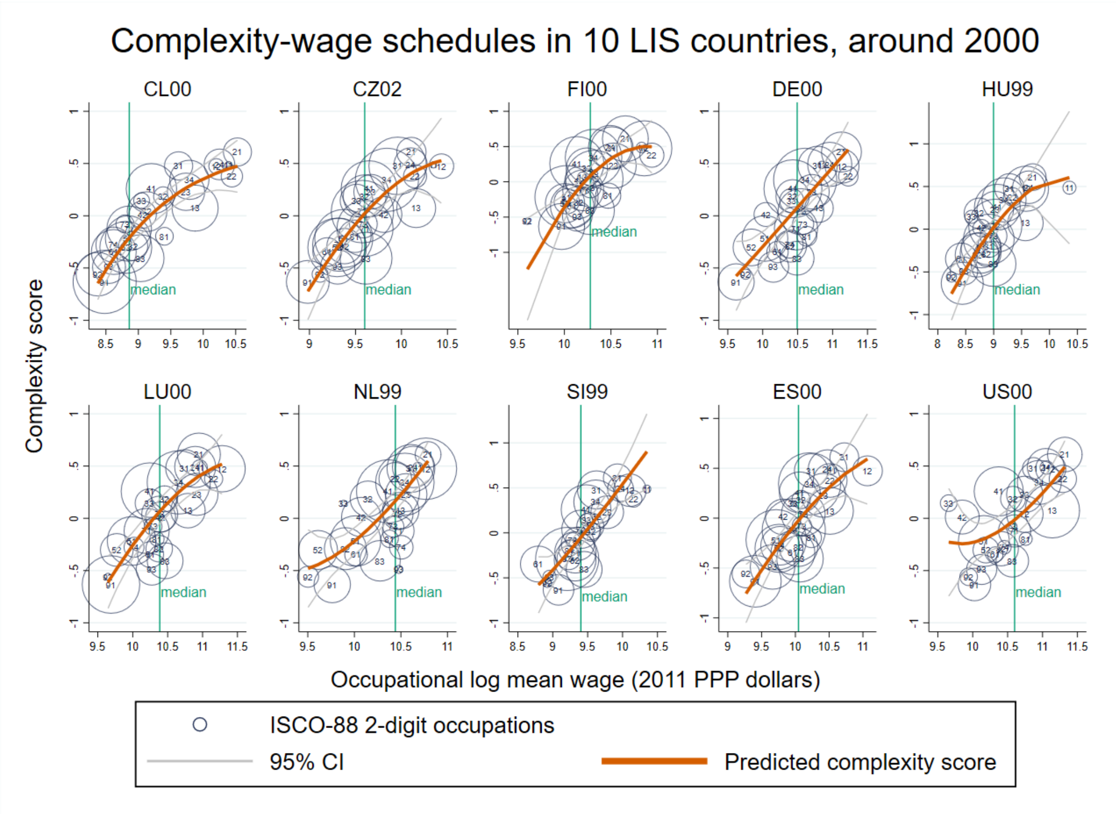

Figure 6.4: *The complexity-wage schedule around 2000. Employment and wage data from LIS wave V (2000 or nearest year).*

Czech Republic and Chile – two countries with historical trajectories very different from those of Germany and the US – the pattern is just as in panel B.1 of figure 6.1. Overall, therefore, it is very clear that there is a straightforward positive relationship between the complexity of the tasks performed and the remuneration received.

6.4.2 The routine-wage schedule

Moving on to the routine-wage schedule, my graphical analyses confirm that there is no universal pattern, in contrast with the complexity-wage schedule. In some countries, routine occupations do cluster predominantly around the middle of the wage distribution as the polarisation argument would predict. However, other countries show the monotonic, linear or near-linear relationship which is associated with occupational upgrading as postulated by Fernández-Macías & Hurley (2017). Figure 6.5 illustrates this for the year 2000. Perhaps most importantly, the textbook examples of employment polarisation, Germany and the US, both exhibit the hump-shaped pattern. This also helps to explain how the idea that routine occupations tend to be medium-wage could gain such widespread acceptance without proper scrutiny: early and influential studies of RBTC often focused on these two countries (see, e.g., Antonczyk, DeLeire & Fitzenberger 2010, Autor et al. 2008, Spitz-Oener 2006). In the US, the hump-shape is more pronounced, and although it is again to a significant extent driven by ISCO group 33, it is robust to excluding associate teaching professionals from the analysis.

A total of four countries exhibit the hump-shaped relationship between RTI and wages: Germany, the United States, Luxembourg, and The Netherlands. This is interesting insofar as those are the four countries with the highest median wages in my sample. On the other hand, countries with lower median wages or where occupational upgrading was particularly pronounced such as the Czech Republic or Chile (presumably due to a catch-up process after the end of the Communist/Pinochet regimes) show a clear quasi-linear monotonic relationship between RTI and wages, as was hypothesised in panel B.2 of figure 6.1. This is *prima facie* evidence to suggest that the level of economic development may play a role in determining the nature of the routine-wage schedule.

The crucial test for my theory is whether the patterns of employment change align with the two clusters of RTI-wage schedules presented here. Indeed, which shape the RTI-wage schedule takes conforms remarkably well to the overall patterns of employment change that I discussed above. The four countries with hump-shaped RTI-wage schedules all experienced employment polarisation (scenario 2). Of the six countries with a (largely) monotonic RTI-wage curve, five experienced occupational upgrading (scenario 1). The exception here is Hungary which has experienced occupational downgrading where upgrading would be expected and thus does not conform to any of the scenarios.[13] This constitutes strong evidence in favour of

[13]The curious case of Hungary can largely be explained by the employment expansion and wage gains in the service occupations of group 51. Personal and protective services workers moved from the 16th to the 61st wage percentile between 1991 and 2015, while expanding their employment share from 2.9 percent to 10.6 percent. This accounts for much of the increase in apparent low-wage employment in figure 6.2. Therefore, the supposed pattern of occupational

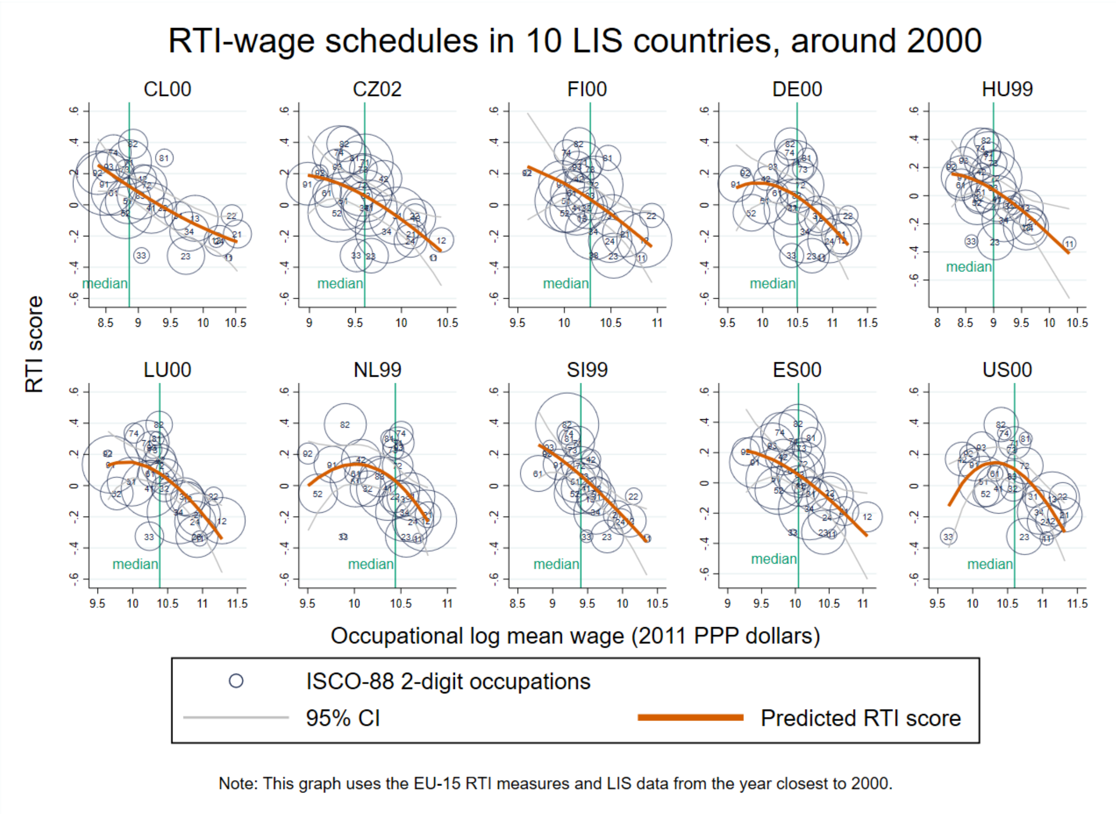

Figure 6.5: *The routine-wage schedule around 2000. Employment and wage data from LIS wave V (2000 or nearest year).*

my theory, as 9 out of 10 cases align with its predictions.

6.4.3 RTI-wage curves and occupational wage rankings over time

It is furthermore important to point out that the shape of the task-wage curves appears to be a relatively constant country-characteristic, even as individual occupations move up or down the wage hierarchy. Figure 6.6 shows the development of the RTI-wage curve in the US, which has the longest time-series in the LIS, from 1974 – 2016. The hump-shaped curve is somewhat less pronounced in later years but is clearly present throughout. Equivalent figures for a further 16 countries are provided in appendix C.3. There are of course minor fluctuations, but instances

downgrading may be better understood as wage upgrading for service occupations. However, if the 1994 – 2013 interval is considered, the changes in group 51 are much less dramatic. The volatility in Hungary of employment shares and relative wages may in turn have to do with the very low sample size compared to all other countries in the analysis.

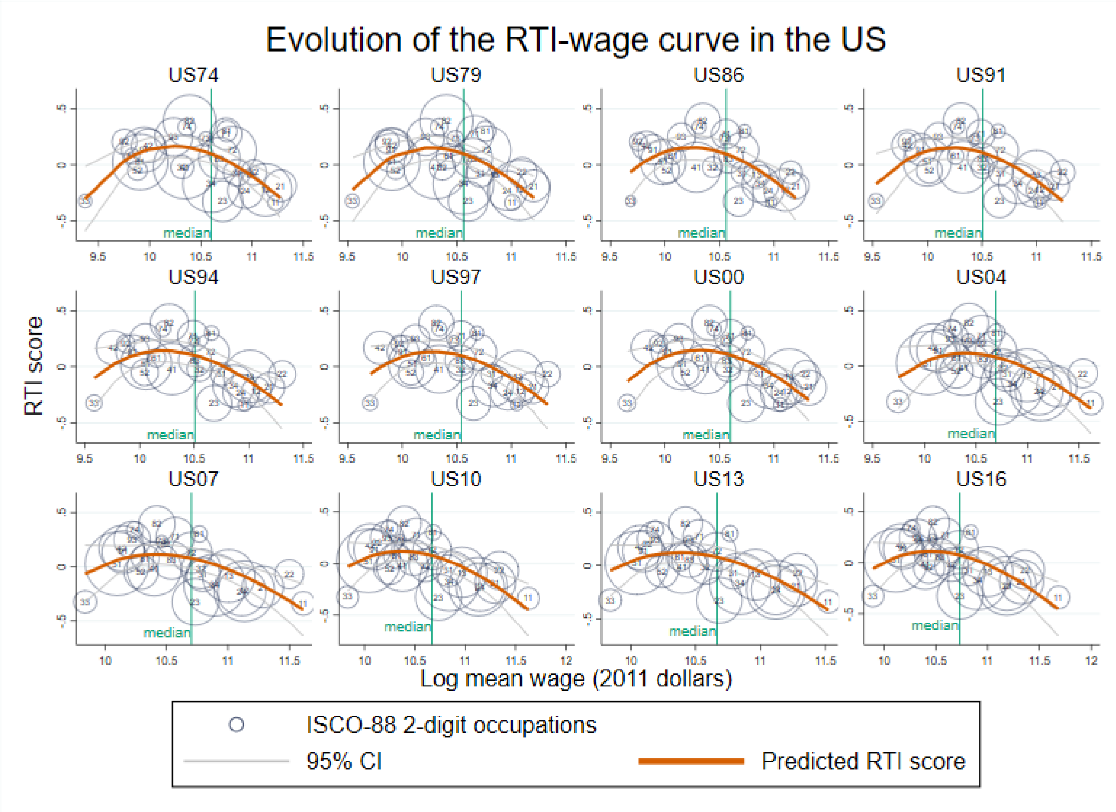

Figure 6.6: *Evolution of the RTI-wage curve in the US from 1974 - 2016.*

where the shape of the RTI-wage curve changes substantially over time are rare. One example is Spain, where a monotonic RTI-wage curve slowly gives way to a marginally hump-shaped one over time. Indeed, where there is sustained and gradual change, as also in Slovenia, it is always towards the hump-shaped relationship. This may indicate that the hump-shaped RTI-wage relationship is characteristic of richer and more technologically advanced countries.

Since the RTI measure does not change over time, all changes in the RTI-wage curve are due to a combination of changing employment shares and occupations moving up or down the wage hierarchy. Substantial employment changes towards higher-paid occupations have already been documented in this study. Large changes in the shape of the RTI-wage curve likely require positional changes in the wage hierarchy as well. In particular, a shift from the linear to the hump-shaped pattern could be the result of the relative wages of routine-intensive occupations increasing over time, or those of a subset of low-to-medium-routine occupations decreasing, or both. The following analyses provide some clarity on this issue.

Table 6.3 shows the three highest-paying and lowest-paying occupations in each country in my sample in 1995 and 2013. We see that the top three occupations are more similar across countries than the bottom three. This is not surprising, as the differences in the RTI-wage distribution that we are trying to explain are at the bottom end. Incidentally, in the countries with hump-shaped RTI-wage curves, service workers (major group 5) appear to be more heavily represented at the bottom than in countries with monotonic curves, where elementary occupations (major group 9) account for most of the three lowest-wage occupations. This is what we would expect, as service occupations are less routine-intensive than elementary occupations. In contrast, there is an overall convergence at the top end of the wage distribution. This indicates that while the top jobs have become more similar across countries, country-specific factors continue to heavily influence the low-wage sector.[14] Thus, despite the relative stability of the RTI-wage curves, occupational wage hierarchies have been far from static.[15] We see in the example of the US that legislators and senior officials (group 11) only become the highest paid occupational group in the early 2000s, after overtaking health professionals (group 22) who had in turn previously overtaken science professionals (group 21) in the early 1990s.

6.4.4 Accounting for non-standard work and standard RTI measures

As mentioned above, I do not exclude non-standard workers, who are known to be disproportionately female, concentrated in certain occupational groups such as service occupations, and to have become more numerous in many countries. For example, Germany and the Netherlands have very high rates of female part-time

[14]Furthermore, the United States' wage structure is exposed as somewhat of an outlier. Nowhere else are teaching associate professionals (group 33) and customer services clerks (group 42) among the lowest earners. As stated above, this does not appear to be driven by part-time employment.

[15]Figure C.9 in appendix C.4 provides a full overview of changes to the wage structure between 1995 and 2013. It shows the average wage percentile of every country-occupation in both years, and the magnitude of changes. This analysis shows that not only low-routine high-wage occupations have generally improved their relative wage position, but also some shrinking high-routine occupations such as operators. The patterns of relative wage changes do therefore not always conform to the predictions of RBTC theory.

Table 6.3: Highest- and lowest-earning occupations by country

Country	1995	2013	Country	1995	2013
Chile	21, 22, 12	21, 22, 12	**Hungary**	11, 12, 13	21, 22, 11
	93, 91, 92	74, 91, 92		93, 92, 91	93, 91, 92
Czechia	12, 11, 13	11, 12, 21	**Luxembourg**	12, 11, 21	12, 24, 11
	93, 92, 91	93, 91, 92		52, 92, 91	52, 92, 91
Germany	11, 12, 21	11, 12, 21	**Netherlands**	11, 12, 21	11, 12, 21
	92, 52, 91	51, 91, 92		51, 52, 91	93, 92, 91
Spain	12, 21, 31	11, 12, 21	**Slovenia**	11, 22, 12	11, 12, 22
	61, 91, 92	51, 91, 92		91, 93, 92	91, 92, 61
Finland	11, 12, 23	11, 12, 22	**US**	22, 21, 12	11, 22, 21
	13, 61, 92	52, 91, 92		92, 42, 33	42, 51, 33

Note: This table shows the three ISCO-88 codes with the highest (top line) and with the lowest (bottom line) average wage for each country in 1995 and 2013. The occupational codes are: 11: legislators and senior officials; 12: corporate managers; 13: general managers; 21: physical, mathematical and engineering science professionals; 22: life science and health professionals; 23: teaching professional; 24: other professionals; 31: physical and engineering science associate professionals; 33: teaching associate professionals; 42: customer service clerks; 51: personal and protective services workers; 52: models, salespersons and demonstrators; 61 market-oriented skilled agricultural and fishery workers; 74: other craft and related trades workers; 91: sales and services elementary occupations; 92: agricultural, fishery and related labourers; 93: labourers in mining, construction, manufacturing and transport. For the full list of ISCO codes, see appendix A.

employment and, together with Finland, the highest overall female employment rates. However, part-time work did not expand uniformly. According to Wright & Dwyer (2003), the share of part-time employment decreased in the US during the long expansion of the 1990s. The OECD (2015) likewise finds decreasing rates of part-time employment in Nordic countries and Spain for the period covered by this study.

Of course, non-standard workers tend to have lower annual labour incomes, all else being equal. Including these workers means that my task-wage curves may be influenced by between-occupation differences and within-occupation changes in patterns of non-standard work. It nevertheless seems to me preferable to include non-standard workers in the analysis. Excluding them would introduce a different kind of bias into the country-comparisons, as non-standard work is an important feature of the labour market, the overall prevalence of which varies between countries and

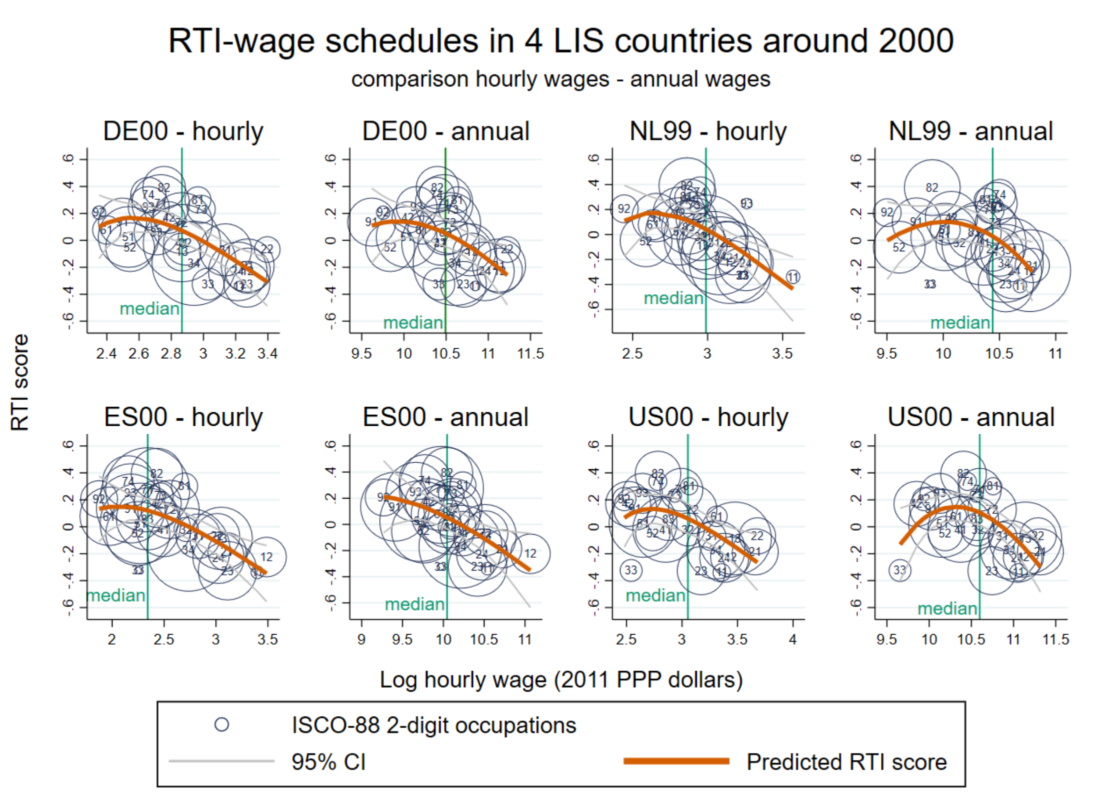

Figure 6.7: *Comparing RTI-wage schedules with hourly and annual wage data, all workers.*

has generally increased. Excluding non-standard workers would mean disregarding an ever-increasing proportion of the workforce, especially women.

Nevertheless, the lack of hourly earnings data constitutes a major disadvantage of the LIS for this type of analysis. Thus, the inability to account for non-standard work due to the use of annual labour earnings invites criticism of my findings. The patterns described so far pertain to all workers and annual wages; hence it is reassuring that figure 6.7 shows that in the four countries with data on gross hourly wages in 2000, the patterns are slightly weaker, especially in the Netherlands and the US, but overall intact.

Appendix C.1 shows further analyses with samples restricted to FYFT workers aged between 16 and 65 and hourly wages. Figures C.4 - C.6 illustrate that the findings presented in figure 6.6 hold true not only in 2000, but also at the beginning and at the end of the period of analysis. Likewise, we see in figure C.7 that of the five countries with information on worker status in 2000, only in Germany do the results for annual wages change when the sample is restricted to prime-age FYFT workers. We can therefore be confident that the results are not predominantly driven by differences in the prevalence of non-standard work, although not in every country the findings are fully robust to using hourly wages or excluding non-standard workers. Nevertheless, the role of non-standard work and gender differences certainly deserves additional scrutiny in future research.

It is furthermore important to note that very similar results obtain when the RTI measures of Autor & Dorn (2013) that Goos et al. (2014) rely on are used. As figure 6.8 shows, the hump-shaped RTI-wage schedule is clearly visible in Germany, Luxembourg, the Netherlands and the United States. In Czechia and Spain, a nascent hump shape is visible while in the remaining countries the monotonic relationship prevails. Thus, even with the more widely known measures of Autor & Dorn (2013), the results described above hold. This shows that the findings are not driven by the choice of routine measure, as the more commonly used routine measure yields very similar results.

6.4.5 Compositional changes within occupations

Another issue is that theoretically, the movements of 2-digit occupations up and down the wage hierarchy that I showed in table 6.3 may represent genuine changes in relative wages as well as compositional changes within the relatively broad 2-digit ISCO categories. It is possible that occupational wages have remained unchanged at the 4-digit level, but, for example, employment growth in higher-earning 4-digit occupations has driven up the average wage of the broader 2-digit group. There are two countries in my sample for which we can assess this possibility in a straightforward manner: Chile and Germany. Both countries report 4-digit ISCO-88 data to the LIS, which have been aggregated to 2-digit level for the main analyses. We can thus investigate the relationship between employment changes at the 4-digit level and the extent to which the wage of a 4-digit occupations deviates from the 2-digit mean. For this analysis I focus on the period 1995 – 2010 (1996 – 2011 for Chile) in order to avoid the confounding influence of the switch from ISCO-88 to ISCO-08.

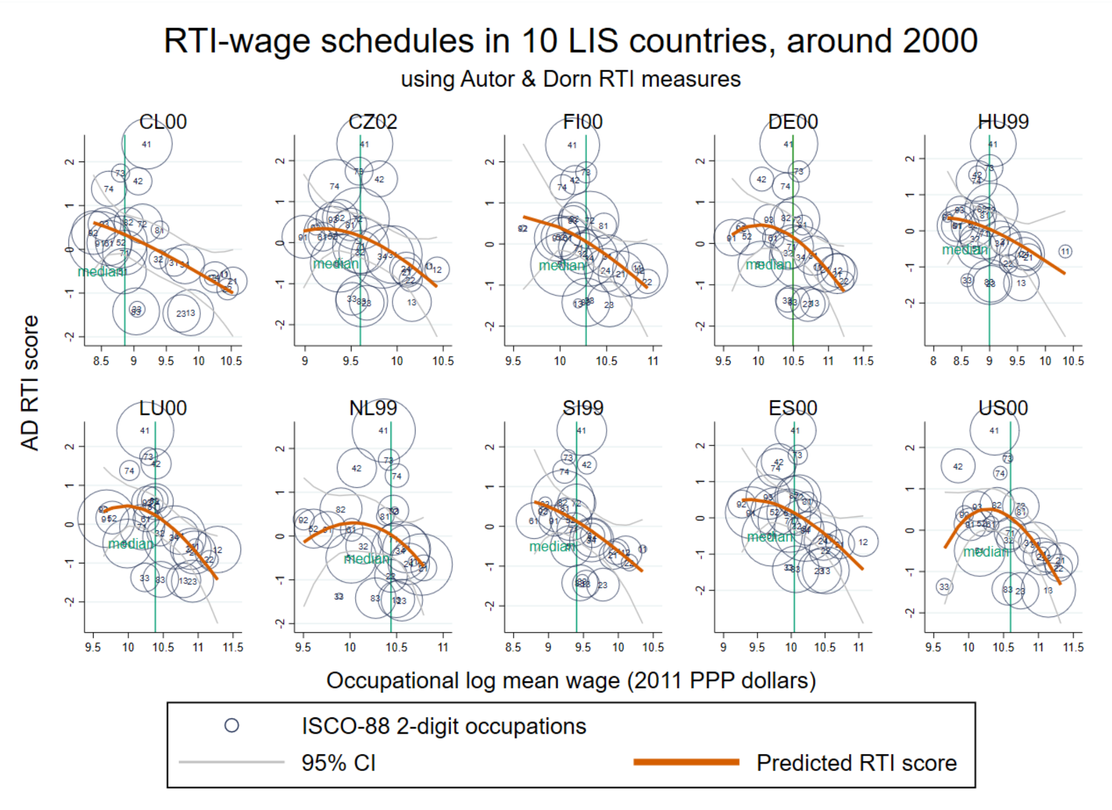

Figure 6.8: *This figure shows that very similar results obtain when the AD RTI measures are used.*

If we plot the change in employment share and the deviation from the 1995 2-digit mean wage of 4-digit occupations in figure 6.9, we see a cross-shaped pattern in both countries. While most occupations, including all large 4-digit occupations, line up on a more or less vertical line around zero deviation from the 2-digit mean wage, there is a less densely populated horizontal line of small 4-digit occupations with more or less constant employment shares and large deviations from the 2-digit average wage. If employment expansion of higher-earning 4-digit occupations was behind the trends at the 2-digit level, we would expect the occupations to be organised around a diagonal line from the bottom left to the top right corner of the graphs. Instead, there is no relationship between employment changes and deviations from the 2-digit mean wage, regardless of whether we look at absolute or relative employment changes and whether or not we control for initial employment.

The impression that compositional changes are not the driving force behind occupational wage trend is reinforced by the analysis of counterfactual wage changes

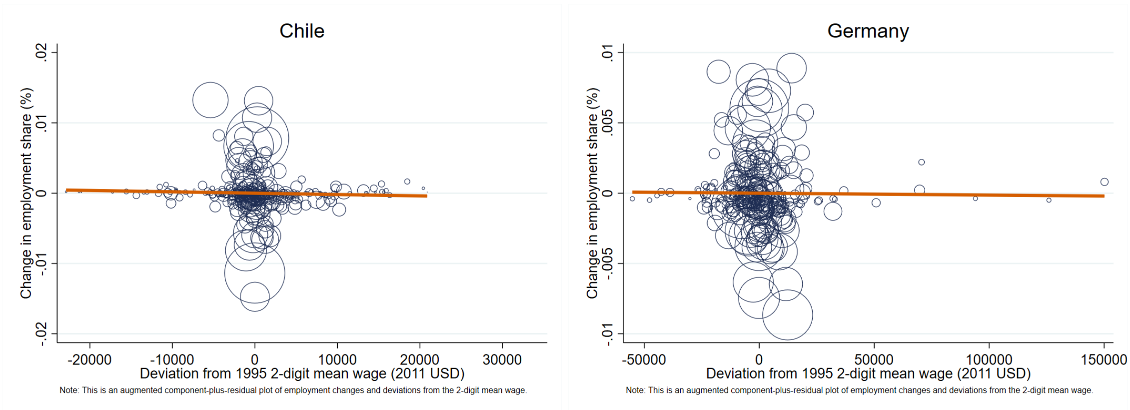

Figure 6.9: *Change in employment share of 4-digit occupation and deviation from mean wage of 2-digit occupation.*

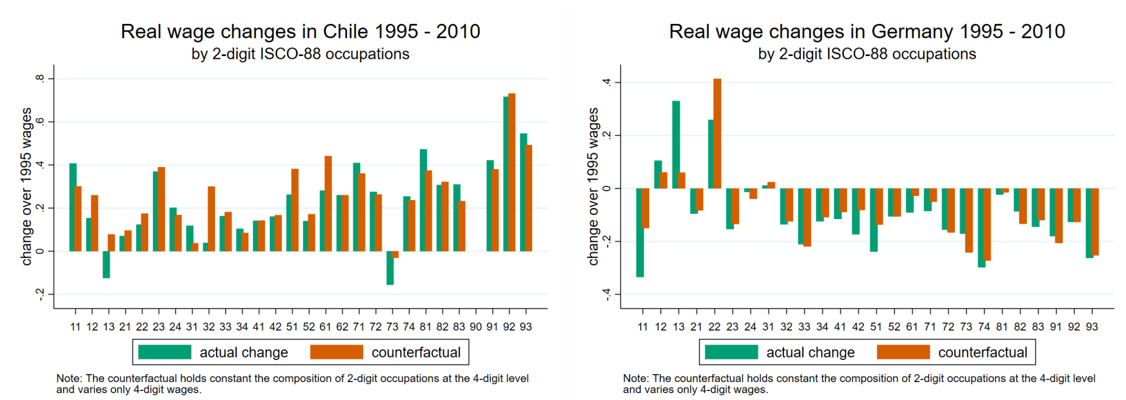

Figure 6.10: *Actual and counterfactual real wage changes at the 2-digit level, holding constant 4-digit employment shares.*

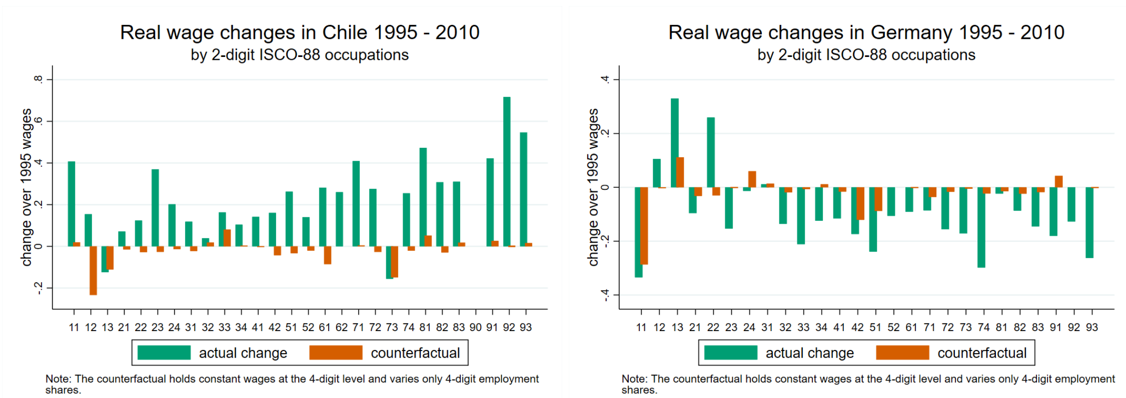

Figure 6.11: *Actual and counterfactual real wage changes at the 2-digit level, holding constant 4-digit wages.*

at the 2-digit level in figure 6.10. For this analysis, I calculated hypothetical 2-digit wages for a scenario in which the shares of 4-digit occupations making up a 2-digit group remained unchanged while 4-digit wages were allowed to develop as they did. As Autor et al. (2008) explain, the validity of this exercise rests on the partial equilibrium assumption that changes in labour market quantities do not affect labour market prices. This runs counter to economic intuitions and findings, but the analysis is nevertheless worthwhile because it can serve as a sanity check. We see that the actual and counterfactual wage changes are very similar for most occupations, regardless of whether we look at Chile which has seen strong overall real wage growth or Germany, where real wage growth has in fact been negative between 1995 and 2010. Repeating the same exercise holding constant wages and only allowing 4-digit employment shares to change, we see in figure 6.11 that the counterfactuals bear little resemblance to the actual changes during this period. This indicates that 2-digit wages would have developed very similarly in the absence of compositional changes at the 4-digit level, and that wage changes at the level of detailed occupations are really behind the majority of wage increases and declines of larger occupational groups.

Overall, the findings for Chile and Germany, while showing that limited compositional changes at the 4-digit level have taken place, do not support the conclusion that these account for the patterns we see at the 2-digit level. Wage changes at the 2-digit level therefore overwhelmingly reflect genuine wage changes, rather than compositional changes, at the 4-digit level. The LIS data do not allow us to perform a similar analysis for every country in the sample, but the very clear results in the case of Chile and Germany should provide some reassurance that compositional changes at the 4-digit level are not a major issue. The findings discussed here provide a starting point for further analyses into the relative importance of wage and compositional changes for explaining changes in the employment structure.

6.5 Conclusion

This chapter takes a critical look at the RBTC argument and proposes a refined theoretical framework to understand the diverse patterns of employment change in countries exposed to very similar technologies. In this framework, an exogenous technological shock leads to an increase in the demand for workers performing complex tasks and a decrease in the demand for workers performing routine tasks. SBTC and RBTC may therefore be at work simultaneously. While task complexity is positively associated with the occupational wage, there are two possible patterns regarding routine-intensity: either routine occupations cluster around the middle of the wage distribution, or the lowest paid occupations are also the most routine-intensive ones. Hence, both occupational upgrading and polarisation may emerge as a result of RBTC. My analyses find empirical support for all components of this refined model, using a novel dataset of 10 countries with employment and wage data from the LIS and task measures based on the EWCS. Thus, the study makes an important theoretical and empirical contribution to the literature on technological and occupational change.

The key empirical finding is that routine occupations either cluster near the middle or near the bottom of the wage distribution in a way that is relatively constant over time but differs systematically between countries. In the former cases, occupational polarisation usually emerges, as posited in the labour economics literature (Acemoglu & Autor 2011, Goos & Manning 2007, Goos et al. 2014). However, in countries where routine occupations earn the lowest wages, occupational upgrading has been the norm. In contrast to other authors who have concluded from this that these countries have experienced SBTC (e.g. Oesch 2013), I argue that their experience is compatible with RBTC if differences in the routine-wage schedules are taken into account.

Furthermore, unlike Fernández-Macías & Hurley (2017) who posit that RBTC in Europe generally results in occupational upgrading and that polarisation, where it does occur, is therefore primarily due to other factors, I show that RBTC can lead to polarisation in the countries where the assumptions of Autor et al. (2003)

and Goos et al. (2014) are met. In future research on technological change and employment dynamics, scholars should take this finding into account – my analyses suggest that one can not simply assume, as many labour economists do, that routine occupations are predominantly medium-wage. Nor are they always low-wage, as some sociologists argue. Hence, employment polarisation is not a necessary consequence of routine-biased technological change, and the upgrading employment changes in many countries are compatible with RBTC, properly understood.

Thus, this study makes several contributions to the scientific debate about how recent technological changes have affected the labour market. It provides a much-needed comparative angle that is not limited to European countries and adds nuance to the debate about the nature of recent technological change. Most importantly, my refined theory can accommodate the different patterns of employment change that have been observed in advanced economies, as RBTC and SBTC are shown to operate simultaneously. Thus, the myth of pervasive polarisation has been further debunked (Oesch & Piccitto 2019), and a plausible explanation for the concomitance of RBTC and occupational upgrading presented.

Currently, the analysis is limited by the nature of the underlying datasets. The LIS, being neither register data nor labour force survey data, does not allow to fully account for part-time work due to missing information on hourly earnings or worker status in many datasets. The smaller sample sizes in some countries also necessitate a higher level of aggregation of the occupational codes, compared to some other studies (e.g. Fernández-Macías & Hurley 2017, Wright & Dwyer 2003). Future research should therefore attempt to replicate the findings in countries for which better data are available.

Nevertheless, the findings in this chapter open up ample opportunities for further investigation. The next logical step in this research agenda is to investigate which structural and institutional factors drive the differences in wage hierarchies which this study uncovered and thereby tighten the link to the extant sociological and economics literatures. Despite the large sociological literature on occupational stratification (see, e.g., Liu & Grusky 2013, Mouw & Kalleberg 2010, Williams 2017), the drivers of differences in occupational wage hierarchies are poorly understood, and in fact, the differences are often overlooked (Hout & DiPrete 2006). Yet, the findings in this chapter suggest that they are crucial for the comparative analysis of employment trends.

Overall, this chapter offers important insights for the study of occupational change and interesting directions for future research. In the next chapter, I investigate some potential drivers of relative wage differences. For example, higher levels of robotisation may contribute to a smaller and better paid routine manufacturing workforce, and thus to employment polarisation. Similarly, labour unions may improve the relative wage position of highly unionised occupational groups such as manufacturing workers. Employment protection for temporary contracts, on the other hand, may reduce the wage gap between low-wage service workers and routine manufacturing workers and thus be conducive to upgrading. By investigating potential heterogeneous effects of robotisation and LMI on the wages of different occupational groups, I integrate the technological and institutional literatures and contribute to a better understanding of the ultimate drivers of employment trends.

7

Wage hierarchies revisited: How robotisation and labour market institutions shape occupational wage hierarchies

Contents

7.1 Introduction

So far in this thesis, I have shown how occupational tasks can be operationalised better and why this matters, and how a refined theory of RBTC centred around the relationship between occupational routine intensity and wages can account for the diverse patterns of employment change in rich countries. Thus, I have provided answers to the first two research questions of the thesis. In this chapter, I follow up on the finding of chapter 6 that routine-wage curves are crucial to understanding the precise implications of RBTC for individual countries. Since the shape of the routine-wage curve is itself an expression of the occupational wage hierarchy, the logical follow-up question is, what determines relative occupational wages?

Therefore, in this chapter, I investigate the third research question that I formulated in chapter 1: do robotisation and LMI, independently and jointly, have heterogeneous effects on the wages of different occupational groups? In doing so, I analyse both the independent effects of robots and LMI and possible moderation effects of LMI on robotisation. The rationale for this is that, to the extent that higher relative wages for routine manufacturing workers are associated with a hump-shaped routine-wage curve, heterogeneous effects of robotisation and LMI on the wages of routine manufacturing and medium- to high-routine service workers would contribute to the emergence of different routine-wage curves.

High-routine manufacturing occupations have experienced variable fortunes with regard to wages, as I showed in the previous chapter. They cluster near the median of the wage structure in some countries and closer to the bottom in others. Moreover, they predominantly employ workers in the industrial sector: 71.7 percent of workers in routine manufacturing occupations work in the industrial sector, compared to 10.7 percent of the remaining workers (see tables 6.1, 7.2, and D.1). The story of varying routine-wage schedules is therefore essentially a story of routine manufacturing occupations.[1] Hence, to understand why the patterns documented in chapter 6 emerge, it is crucial to understand why routine manufacturing occupations are relatively better paid in some countries than in others. To this end, I investigate the manufacturing wage premium, the ratio of the average log wage of routine manufacturing occupations and the average log wage of non-manufacturing medium- to high-routine occupations. In 2010, this premium was 25 percent in the polarising countries analysed in chapter 6, but only 6 percent in the upgrading ones (see figure 7.4).

[1]There are of course routine-intensive non-manufacturing occupations, such as agricultural labourers and office clerks. However, due to the concentration of routine work in manufacturing occupations and the interest in robotisation, it is sensible to focus the analysis on routine manufacturing occupations.

I pursue three related arguments: The first, productivity argument is based on the labour economics literature and holds that the manufacturing wage premium is a function of relative labour productivity in routine manufacturing and non-manufacturing occupations. Robots are assumed to increase the productivity of the workers they are employed alongside, and thus the relative productivity of the manufacturing sector vis-à-vis non-manufacturing. Moreover, countries which are closer to the technological frontier undergo a given technological change earlier than other countries. Hence in these countries the automation of routine employment has begun earlier, and the remaining routine workers are relatively more productive and should command higher relative wages (Graetz & Michaels 2018). This logic relies on the economic argument that ultimately, workers' wages depend on their productivity, and a technology-induced increase in productivity will therefore raise the wages of the affected workers (Autor 2014).

The second, institutional argument draws on power resource theory and a substantial literature in economic sociology and comparative politics on the institutional determinants of wages and inequality (Korpi 1983, Rueda & Pontusson 2000, Western & Rosenfeld 2011), as well as an emerging literature on the political consequences of automation (see, e.g., Thewissen & Rueda 2017). What is missing in this existing literature is an analysis of how labour market institutions structure the wage hierarchy through their differential impact on various occupational groups. Based on the logic of power resource theory, occupational groups that have greater exposure to inclusive labour market institutions such as unions and employment protection improve their relative bargaining position which should be reflected in relative wages.

The productivity and power resource arguments are powerful explanations in their own right and as such have so far not been studied jointly. Therefore, the third argument I investigate is that LMI may moderate the impact of robotisation on wages. The literature suggests that workers will use their bargaining power to shape the adoption of new technology in their interest (see, e.g., Fernandez 2001). Thus, for example, in the presence of strong unions, robots may be associated with a higher manufacturing wage premium, as unions negotiate retraining opportunities for workers to improve their productivity (Fernandez 2001). Alternatively, if unions oppose robotisation, higher robot density may indicate eroding union power and thus be associated with a lower wage premium (Parolin 2021). To my knowledge, this is the first cross-country study to explicitly test the hypothesis that LMI moderate the impact of robotisation on wages.

I find evidence in favour of the productivity and institutional arguments in descriptive analyses; however, only the institutional argument with regard to employment protection holds up in a more thorough econometric analysis. Robotisation and union density appear to be associated with a lower manufacturing wage premium. Moreover, I find evidence that employment protection, although not union density, moderates the relationship between robotisation and the manufacturing wage premium. This bolsters the argument made in chapters 2 and 3 of this thesis that those who disregard institutional factors in an economy seemingly driven by technological change, are missing an important part of the puzzle.

This chapter thus follows up on my previous findings and offers an explanation for what causes the differences in occupational wage hierarchies that can have such far-reaching consequences. The approach pursued here is innovative in two main ways. The first novelty in my analyses is the use of a theoretically motivated wage ratio as the dependent variable. To my knowledge, no previous study has specifically analysed the impact of robotisation on the relative wages of manufacturing workers who are the primary target of this type of automation. Moreover, for all that has been written about the impact of labour market institutions on overall inequality, little is known about which particular occupational groups are the main beneficiaries of inclusive institutions. The second innovation is the explicit consideration of a potential moderating effect of LMI on the relationship between robotisation and relative wages. Despite the obvious potential for further integration of the economics and sociological literatures, such efforts have so far been wanting. My analysis addresses this gap.

The chapter proceeds as follows. Section 7.2 reviews the literature on wage hierarchies and the impact of robots and labour market institutions on relative wages.

Section 7.3 sets out a theory of occupational wage hierarchies and formulates a set of hypotheses. Section 7.4 introduces the data and the analytical strategy. Section 7.5 presents descriptive evidence and section 7.6 the main empirical findings. Section 7.7 concludes the chapter with a discussion of the findings and their implications.

7.2 Literature overview

In this section, I discuss a range of existing research that has not yet been covered extensively in chapter 2. For the agenda of this study, three important issues must be addressed. First, my argument about different occupational wage hierarchies is put in the context of the sociological literature on the uniformity of occupational hierarchies. Second, with a view to my productivity argument, the emerging literature on the labour market impact of robots merits a detailed discussion. Finally, I recapitulate what the literature has to say about the distributional impacts of LMI at the occupational level.

7.2.1 How much uniformity in occupational wage hierarchies?

Despite sociologists' longstanding interest in social inequalities, the analysis of wage hierarchies has been comparatively neglected. In some cases, how much workers earn in one occupation relative to others may differ substantially between countries – just consider the prominent case of teacher pay: teachers are paid at the 78th percentile of the wage distribution in Korea and at the 49th percentile in the United States (Dolton & Marcenaro-Gutierrez 2011). That wage hierarchies have not attracted more interest may be because, despite these examples, by and large, wage hierarchies are very similar.

In fact, Hout & DiPrete (2006, p. 2) list the finding that "occupations are ranked in the same order in most nations and over time" as the first of 19 empirical generalisations about social stratification that sociologists have established. This so-called "Treiman constant" dates back to Don Treiman's pathbreaking cross-country study of occupational prestige which established this pattern and has been confirmed many times since (Treiman 1977).[2] Hout & DiPrete (2006, p. 3) refer to it as perhaps "the only universal sociologists have discovered – not just in stratification but sociology as a whole." Similarly, Avent-Holt et al. (2020, p. 1) call it "one of the most reliable findings in the social sciences."

Based on data from 85 studies of 60 societies, Treiman (1977) found an average correlation across countries of 0.79 based on the Standard International Occupational Prestige Scale (SIOPS), with a standard deviation of 0.14. This is remarkable; yet, the slight variations around the overall hierarchy established by Treiman suggest that there may be country differences that are the result of other labour market characteristics and which in turn affect employment outcomes. The results in chapter 6 indicate that this is indeed the case.

Regarding deviations from Treiman's hierarchy, Avent-Holt et al. (2020) note that less developed regions and the former socialist countries correlate less well to the scale, as they for example rate manual occupations higher than industrialised Western countries. They argue that such variations are important, but "should be understood as variations on an underlying uniform occupational hierarchy" (Avent-Holt et al. 2020, p. 2). le Grand & Tåhlin (2013) reiterate the largely invariant rank order of occupations in 11 European countries but point out the great differences in wage differentials between the ranks. They moreover argue that efficiency mechanisms, rather than the power mechanisms emphasised in theories of social class, determine the rank order of occupations.

Thus, there appears to be widespread agreement that wage hierarchies are quite uniform, but what gives rise to this uniformity – and any deviations from it – is more contested and, above all, less well documented. Moreover, the existing literature, even as it emphasises uniformity, does not dismiss the possibility of

[2] While the "Treiman constant" refers to occupational prestige, rankings based on different measures of occupational quality produce virtually identical rankings. For example, Ganzeboom and Treiman's (1996) rankings based on the SIOPS and International Socio-Economic Index of Occupational Status (ISEI) have a rank order correlation at the 2-digit ISCO-88 level of 0.86. Avent-Holt, Hällsten & Cort (2020, p. 2) "refer to the Treiman constant as an empirical phenomenon documenting a roughly invariant occupational hierarchy, regardless of how this occupational hierarchy is measured or conceptualized." They then go on to use wages to operationalise occupational hierarchies in Swedish workplaces.

small but potentially important differences manifesting themselves. The findings of chapter 6 provided a crucial piece of evidence that such differences are real and empirically meaningful. The analyses in this chapter seek to determine why they emerge in the first place.

7.2.2 Technological change, robotisation, and the labour market

How to best operationalise technological change has been a longstanding problem in the technological change literature. Often, it has simply been treated as a "black box". Researchers who analyse technological change as a demand side phenomenon tend to use it as a "residual concept, whose operational meaning is often 'labor demand shifts with invisible causes'" (Bresnahan, Brynjolfsson & Hitt 2002, p. 340). Just how technology changes the demand for different types of labour has been notoriously difficult to pin down. Of course, the notion that RBTC affects the demand for certain kinds of tasks represented a major advance in this area (Autor et al. 2003). Yet, it still did not solve the fundamental issue of attribution, as all manner of other factors may also influence task demands.

Commonly used direct measures of technological change include productivity improvements or declines in the price of computer capital. For example, Autor et al. (2003) model the declining price of computer capital as the main factor driving the reallocation of routine tasks to computers and machines. In addition to these macro-level measures, in recent years, data on industrial robots have increasingly become available. This has made it possible to operationalise automation as one crucial manifestation of technological change in a more straightforward manner. The mechanisms underlying RBTC theory can so be investigated more directly, although one has to keep in mind that robotisation is not an encompassing measure of RBTC. This literature constitutes the first pillar of my argument in this chapter.

As the study of robots is still in its infancy, few consensus findings exist on any of the key issues, as is evident in table 7.1 which summarises the scope conditions and findings of the most important studies currently available on the topic of robotisation.[3]

Robots and employment

Regarding employment, the results seem to suggest that robot use is linked with higher employment at the level of individual firms (Acemoglu, LeLarge & Restrepo 2020). This may however be due to a capture of market share from less automated and therefore less productive firms. Consistent with this interpretation, at the regional and sectoral level, negative employment effects on low-skill workers, the manufacturing sector, or in more heavily exposed commuting zones appear to prevail (Acemoglu & Restrepo 2020*a*, Dauth et al. 2019, Graetz & Michaels 2018). This constitutes evidence for a displacement effect of robotisation.

Yet, these trends are to some extent counteracted by reallocation effects, meaning for example that job losses in manufacturing due to robotisation are offset by employment growth in local services (Dauth et al. 2019). The impact on aggregate employment is therefore often judged to be zero. Dauth et al. (2019) estimate, for example, that while each new robot installed in Germany has on average destroyed two manufacturing jobs, employment gains outside manufacturing entirely compensated for this decline. Graetz & Michaels (2018) and Klenert et al. (2020) find similar reallocation effects. This is also in line with theoretical work that emphasises how technology not only displaces but also can reinstate labour (Acemoglu & Restrepo 2019, 2020*b*, Autor & Salomons 2018).

The employment effects of robotisation do not affect all demographic groups equally. Chiacchio, Petropoulos & Pichler (2018) point out that negative employment effects of robotisation in Europe tend to fall disproportionately on groups such as younger workers and male workers. Acemoglu & Restrepo (2020*a*) for the US likewise find a stronger negative employment effect on men which is concentrated

[3] A robot is defined as a programmable mechanism with a degree of autonomy that can move within its environment to perform intended tasks (International Organization for Standardization 2012). Klenert et al. (2020) report that industrial robots in Europe overwhelmingly perform physical tasks such as handling operations and machine tending or welding and soldering, which are concentrated in the manufacturing sector.

in manufacturing, whereas female employment declined more in non-manufacturing. Dauth et al. (2019) show that in Germany young labour market entrants bear the brunt of the decline of manufacturing jobs and adjust by taking service jobs.

The verdict is therefore still out to what extent and how robotisation has contributed to the decline of routine employment. As Klenert et al. (2020) point out, the decline of manufacturing employment precedes the advent of robotisation and does not appear to have been accelerated by it. It is therefore not the sole, or even the main reason for the decline. Indeed, Kollmeyer (2009) investigates the causes of deindustrialisation in 18 OECD countries from 1970 until 2003 and argues that direct and indirect effects of globalisation, growing affluence, and disproportionate productivity growth have all contributed to reduced manufacturing employment. Against this background, Klenert et al. (2020) find that across Europe between 1995 and 2015, country-sector pairs with high levels of robot adoption have been more resilient to the decline of manufacturing employment. The results from these studies indicate that robots may not be directly destroying jobs – routine or otherwise – at a large scale, contrary to the popular perception (Mishel & Bivens 2017).

Robots and wages

More directly relevant for this study is the impact of robotisation on the wages of routine workers. Conventional theory holds that as RBTC automates routine jobs, it should also depress their wages (Acemoglu & Autor 2011). The cross-country evidence for this, however, is mixed at best, as chapter 2 has shown (see, e.g., Antonczyk et al. 2018, Michaels et al. 2014, Naticchioni et al. 2014). Specifically with regard to robots, there is as yet little consistent evidence of an effect in any direction.

Graetz & Michaels (2018) to date is the only major cross-sectional study of robots and the labour market that considers an effect on wages. Their sample includes sector-specific data from 17 countries and their regressions find a positive effect of robot adoption on changes in mean hourly wages. However, this effect is approximately an order of magnitude smaller than the effect on total factor productivity, suggesting that gains from improved productivity are only partly passed on to workers. Particularly interesting is their finding of heterogeneous effects by skill group: robotisation appears to reduce the share of hours worked (and thereby annual or weekly wages, *ceteris paribus*) of low-skilled workers, and has a small positive effect on the hours worked of middle- and high-skilled workers. As the authors note, this is noteworthy given the finding in the RBTC literature that technological change is biased against medium-skill workers. It is, however, completely in line with the argument investigated in this study, that medium-skill workers can be beneficiaries of robotisation in terms of wages. Since this study looks at annual rather than hourly wages, an increase in the share of hours worked translates directly into higher relative wages.

Chiacchio et al. (2018) perform a similar region-sector analysis on a more limited scale for 6 European countries. They find no consistent wage effect of changes in robot exposure, although the results lean towards a negative relationship. The authors speculate that the insignificant overall relationship may be the result of counterbalancing developments at the firm level, such as investment in human capital. While macro-evidence for this is hard to come by, it is for example what Fernandez (2001) found in his study of a plant retooling. Interestingly, the authors show that nominal gross hourly wage growth for plant and machine operators and assemblers in the 1995 – 2007 timeframe has been among the highest of all occupational groups, even though they are heavily affected by technological change. They find no evidence, however, that this increase is due to robotisation. Their findings therefore do not allow to form any strong conclusions.

Moving on to single-country studies, Acemoglu & Restrepo (2020*a*) estimate wage effects across the wage distribution in the US and find a significant negative effect of greater robot exposure in the bottom third of the distribution. For non-college educated workers, the estimated effect is negative up until the 85th percentile, whereas for the college-educated, negative effects concentrate below the 15th percentile. Looking at women and men separately, the effects for women in the bottom half of the distribution lean negative but are not consistently significant, whereas for low-wage men, there is a consistent negative effect up until approximately the median wage.

Dauth et al. (2019) report that the adoption of industrial robots by German firms reduced overall manufacturing employment by reducing the number of young entrants, while incumbent workers often took on new roles within the firm that are of higher quality and earn higher wages. Thus, the wage effect of robotisation is positive for workers who stay in the same firm, but negative for those who change employers. This appears potentially consistent with a positive wage effect of robotisation for the remaining manufacturing workers, depending on the share of workers who separate. However, in a similar analysis Acemoglu et al. (2020) found no consistent wage effects for robot-adopting firms in France.

It is therefore clear that a lot remains to be learned about the impact of robotisation on wages. Cross-country evidence is scarce and the scalability of findings of single-country studies is questionable. Furthermore, the impact on the relative wages of different groups is not well understood. Graetz and Michaels's (2018) and Dauth et al.'s (2019) studies suggest that workers in occupations and sectors that experience robotisation may benefit in terms of their wages, while Chiacchio et al. (2018) and Acemoglu et al. (2020) report inconclusive findings. According to Acemoglu & Restrepo (2020*a*), negative wage effects dominate in the US. Yet, the recently available robot data afford an opportunity to investigate the role of robots in shaping the occupational wage hierarchy in greater detail. By doing so, we moreover stand to gain insights into the mechanisms that are at the heart of RBTC. However, it is important to remember that industrial robots are just one facet of technological change, and findings with regard to robotisation cannot necessarily be extrapolated to other manifestations of RBTC.

Table 7.1: Overview of studies of robotisation

Study	Countries	Time	Level of analysis	Employment	Wages	Productivity
Chiacchio et al. 2018	6 EU countries	1995 – 2007	Region-sector	Negative effect, especially for middle educated, young, male workers	No consistent wage effects (leaning negative)	N/A
Graetz & Michaels 2018	17 advanced countries	1993 – 2007	Country-sector	Reduce low-skill employment, not total employment	Small increase in overall wages	Increase labor productivity growth by 0.36%
Dauth et al. 2019	Germany	1994 – 2014	County/ firm	Displacement in manufacturing offset by new service jobs	Wage increases in robot-adopting firms, overall effects negative	Robots associated with significant increase in labour productivity
Acemoglu et al. 2020	France	2010 – 2015	Firm	Robot-adopting firms expand employment; overall impact negative	No consistent wage effects	Positive effect on labour productivity
Acemoglu & Restrepo 2020*a*	US	1990 – 2007	Commuting zone/ sector	Reduction of overall employment, especially routine occupations	Reduction of overall wages, especially low-wage workers	Positive effect on labour productivity
Cséfalvay 2020	18 EU countries	1995 – 2015	Country-sector	N/A, looks at robot growth instead	N/A	N/A
Klenert et al. 2020	EU-28 countries	1995 - 2015	Country-sector	Robot use linked to increase in aggregate employment, at least in manufacturing	N/A	N/A

7.2.3 Labour market institutions and relative occupational wages

As I argued in chapter 2.2, a major shortcoming of the SBTC and RBTC literatures is that they have developed with very little reference to a rich body of literature on the impact of labour market institutions on the wage distribution, including in economics (see, e.g., Card & DiNardo 2002, Card, Lemieux & Riddell 2003, Oesch 2013). This literature is based on power resource theory which posits that labour market institutions such as collective bargaining, employment protection, or a minimum wage equip workers with greater bargaining power in the struggle over the distribution of resources (Esping-Andersen 1990, Korpi 1983). Generally, this is associated with a compressed earnings distribution, especially at the bottom (Bryson, Ebbinghaus & Visser 2011, Koeniger et al. 2007, Rueda & Pontusson 2000).

Yet, labour market institutions do not affect all workers and occupations equally (Eichhorst & Marx 2015). This is obvious, but the implications of it have not been studied comprehensively. For example, it is well established that union-covered workers earn more, on average, than non-covered workers (Breda 2015, Freeman & Medoff 1984, Visser & Checchi 2011). Parolin (2021) shows for the US that organised labour still plays an important role under conditions of technological change, protecting wages in highly unionised sectors even at the expense of a faster employment decline. Comparative analyses of the impact of different levels of unionisation on relative occupational wages, however, are lacking.

Traditionally, unionisation has been highest in the manufacturing sector in many countries (Card et al. 2004, DiNardo et al. 1996, Ebbinghaus & Visser 2000). This would boost the relative wages of routine occupations which tend to be in the manufacturing sector. In Germany, for example, the manufacturing sector boasts a high union density and manufacturing wages are high compared to those in the service sector (Dribbusch, Lehndorff & Schulten 2018).[4] To the extent that manufacturing workers are indeed more likely to be unionised, extrapolating from

[4] Of course, the German manufacturing sector, with its heavy reliance on premium car manufacturing and high-tech machinery is highly productive compared to manufacturing in many other countries.

established findings suggests that unionisation may boost the relative wages of routine manufacturing workers, but there is as yet no empirical evidence.

Employment protection legislation (EPL) increases worker bargaining power by making dismissals costly for employers. Its core function is thus to reduce labour turnover, and it has been described as an “employer-borne tax on employment adjustment” (OECD 2004, p. 65). EPL has also been linked to lower wage inequality and employment and higher unemployment (Koeniger et al. 2007, MacLeod 2011, Pissarides 2001, Salverda & Checchi 2015), but there is so far no direct evidence for an effect on occupational wage hierarchies.

The EPL indicator that was developed by the OECD and is used in most empirical analyses distinguishes between permanent contracts and temporary contracts. Most empirical studies focus on the effect of regulation of permanent contracts on employment, turnover, and wages, as this is the dominant mode of employment (Koeniger et al. 2007, Nunziata 2005). This is unfortunate, since much of the reform activity in recent decades relates to temporary contracts, and as restrictions on temporary contracts have been relaxed, non-standard employment has become more prevalent in many countries (Weisstanner 2020).

Yet, the prevalence of temporary employment differs substantially between countries and between sectors (Eichhorst & Marx 2015). For example, temporary employment is more common in the service sector than in the manufacturing sector. Thus, even more clearly than with unions, not all occupational groups are equally affected by EPL for temporary workers. A study by Weisstanner (2020) finds that the deregulation of non-standard employment has adverse wage effects even on low- and middle-income workers on standard contracts in a sample of 22 democracies. This suggests that the wage effects of EPL for temporary workers may be substantial, as not only non-standard workers are affected. Evidence for spillover effects of employment deregulation is also provided by Arestis et al. (2020). More research on the wage effects of EPL for temporary contracts is sorely needed, as there are few other studies investigating this link. Like above, assuming that EPL works analogously for relative occupational wages, one may expect a negative relationship between EPL for temporary contracts and the wage premium in manufacturing occupations, but direct evidence is lacking.

So far, moreover, the impacts of robotisation and labour market institutions on occupational wage hierarchies have been treated separately by most scholars. Yet, it is increasingly recognised that they do not exist in a vacuum. For example, Chiacchio et al. (2018, p. 23) speculate that "European labour market policies could possibly cushion the impact of industrial robots, leading to a less severe drop in the employment rate", and call for more research in this area. Aghion, Antonin & Bunel (2019) also hypothesise, based on an analysis of robotisation in France between 1994 and 2014, that labour market institutions and policies broadly conceived may significantly influence the effect of robotisation on employment. In particular, they point to the importance of education policy in mediating the labour market impact of robots by creating complementarities. Dauth et al. (2019) also find evidence for a role of labour market institutions in explaining differences in the displacement effect of robotisation.

A more forceful line in favour of institutions is taken by Mishel & Bivens (2017) who argue that robotisation and automation cannot explain recent increases in inequality and do not lead to overall employment reductions. While only Mishel & Bivens (2017) directly investigate the relationship between automation and institutional factors, and none of these contributions investigate specific implications for occupational groups such as routine workers, this overview makes it clear that in recent years, it is increasingly recognised that automation and institutional factors should not be treated in isolation.

This literature overview has shown that while occupational wage hierarchies are generally assumed to be invariant across countries, this does not preclude small but meaningful variations, such as I identified in chapter 6. However, research on the impact of technological and institutional factors on differences in the relative wages of different occupational groups is thus far almost non-existent. In this chapter, I extend existing research by addressing this gap.

7.3 A theory of occupational wage hierarchies

7.3.1 Theoretical framework

In this chapter, I draw on the RBTC and robotisation literature as well as on power resource theory and the comparative institutions literature to investigate whether robotisation and labour market institutions independently and jointly influence the relative position of routine manufacturing occupations in the wage distribution. Before formulating the hypotheses, it is helpful to review again the different routine-wage curves that were described in chapter 6 (see figure 6.5). Panel 1 of figure 7.1 represents the reality for most countries: a monotonic RTI-wage curve which suggests a relationship that is strictly negative but not exactly linear. Panel 2 shows a hump-shaped curve as it is present in countries like the US. The curve exhibits a local maximum in the lower half of the wage range, meaning that high-routine occupations are concentrated just below the median wage. In chapter 6, I have found that polarisation is limited to countries with the hump-shaped routine-wage curve, while the remaining countries overwhelmingly experienced occupational upgrading. The crucial difference between the two patterns is in how the average wage of high-routine occupations compares to that of medium-routine occupations (low-routine occupations have the highest average wages in all cases). To analyse the impact of robotisation and LMI on wage hierarchies and, by extension, employment change, this chapter therefore looks at a high-routine/medium-routine wage ratio, as is explained in detail in the next section.

7.3.2 The manufacturing wage premium

A key conceptual challenge for this study is the operationalisation of a suitable dependent variable. Since the research question refers to the relationship between groups of occupations, an aggregate measure at the country-year level is necessary. In principle, I am interested in the ratio of average wages in high-routine and medium-routine occupations. However, since I also analyse the impact of robotisation and almost all robots are employed in the industrial sector (Klenert et al. 2020), it

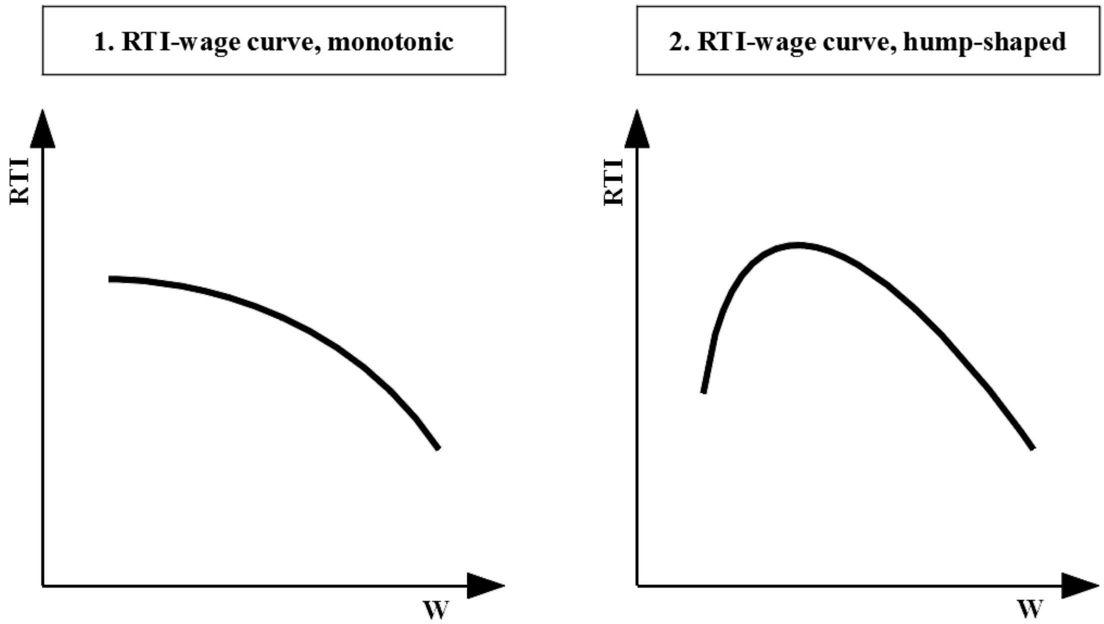

Figure 7.1: *Different shapes of the RTI-wage relationship: monotonic and hump-shaped RTI-wage curves.*

makes sense to focus specifically on high-routine manufacturing occupations. To this end, I calculate a ratio which I call the manufacturing wage premium.

The construction of the manufacturing wage premium is based on a few simple rules. All 2-digit occupations that are

1. among the most routine-intensive third of employment-weighted occupations based on the measures from chapter 5, and
2. in which a majority of workers are employed in the industrial sector according to table D.1,[5]

are classified as high-routine manufacturing occupations and constitute the numerator of the manufacturing wage premium. Conversely, all occupations that are

1. not among the least routine-intensive third of employment-weighted occupations, and

[5]This analysis is based on data from the EWCS for the EU-15 countries, since with the occupational recodes the LIS data could only be used for the countries which originally report ISCO-88 data. However, the allocation of occupations to sectors is very similar across countries and it is highly unlikely that there would be meaningful differences between the samples.

2. in which a majority of workers are not employed in the industrial sector,

are designated as medium- to high-routine non-manufacturing occupations and make up the denominator of the manufacturing wage premium.

Based on this allocation, I define P_{cy} as the the average log wage of high-routine manufacturing occupations HRW_{cy} divided by the average log wage of medium- to high-routine non-manufacturing occupations MRW_{cy}. Thus, the manufacturing wage premium can be expressed as:

$$P_{cy} = \frac{\ln\left(HRW_{cy}\right)}{\ln\left(MRW_{cy}\right)}. \tag{7.1}$$

Table 7.2 lists the occupational groups that are part of either group; the remaining, mainly professional occupations, make up the low-routine, high-wage group which is of limited interest for my research question in this chapter.

With this definition of the manufacturing wage premium, I focus on the lower two thirds of the occupational hierarchy. The occupations in the numerator on average account for 25 percent of total employment in the sample used in the regression analyses below, while the occupations in the denominator make up 41 percent. This region is where most of the differences between countries are located (see chapter 6), and where the effects of robotisation and institutions are expected to play out (Fernández-Macías 2012).

With regard to robotisation, the crucial assumption from RBTC theory is that it predominantly affects high-routine occupations. While some service occupations are also routine-intensive, it is generally speaking software and not robots that may substitute for routine tasks in these occupations. Effects from robotisation should thus be concentrated in the occupations that make up the numerator of the manufacturing wage premium.[6]

Regarding the institutional argument, Fernández-Macías (2012, p. 176) states, “the fact that most of the differences between countries lie in the bottom quintiles

[6] As the robot data are sectoral and my wage data occupational, some measurement error is inevitable. This is why my allocation rules cross-classify occupations by their dominant sector of employment based on table D.1.

Table 7.2: Occupational groups comprising the main dependent variable

High-routine manufacturing		**Medium- to high-routine non-manufacturing**	
ISCO	**Occupation name**	**ISCO**	**Occupation name**
71	Extraction and building trades workers	**32**	Life science and health associate professionals
72	Metal, machinery and related trades workers	**41**	Office clerks
73	Precision, handicraft, craft printing and related trades workers	**42**	Customer service clerks
74	Other craft and related trades workers	**51**	Personal and protective services workers
81	Stationary plant operators	**52**	Models, salespersons and demonstrators
82	Machine operators and assemblers	**61**	Skilled agricultural and fishery workers
93	Labourers in mining, construction, manufacturing and transport	**83**	Drivers and mobile plant operators
		91	Sales and services elementary occupations
		92	Agricultural, fishery and related labourers

provides … support for the institutional argument." Unionisation and EPL are indeed widely acknowledged to exert their influence mainly in the lower parts of the wage distribution, and the discussion above suggests that there may be differences between high-routine manufacturing and medium- to high-routine non-manufacturing occupations which this analysis seeks to exploit (Koeniger et al. 2007, Rueda & Pontusson 2000).

Thus, the manufacturing wage premium is operationalised so as to obtain maximum leverage for both explanatory approaches. One could think of it as a theoretically informed alternative to the widely used 50/10 wage ratio.[7] Instead

[7] Cortes (2016) also investigates wage premia for groups of occupations, but in his operationalisation, the industry angle which is crucial for my analysis of robotisation is not considered.

of measuring the impact of, say, unions on inequality by investigating percentiles of the overall wage distribution, this approach measures the relative wages of two groups of occupations that theory suggests should be exposed to unions to different degrees. The relative wage ratios thereby capture numerically what the routine-wage curves show graphically, and can serve as the dependent variable in the statistical analyses in this study.

7.3.3 Hypotheses

Industrial robots are a manifestation of technological change, and higher robot density suggests greater technological advancement. The first central hypothesis of this study is that while robotisation may reduce overall employment in high-routine manufacturing occupations,[8] robots complement the remaining routine workers and raise their relative wages.[9]

Complementarity in this context manifests itself in productivity-enhancing upskilling: a worker in one of the high-routine occupational groups (such as plant and machine operators) who previously handled machines on the assembly line may see his function change to supervising an automated assembly line staffed by robots. The broad occupational classification would be the same, many of the tasks may also be very similar, but the added value and hence the wage of the worker would be significantly higher.[10] There is micro-evidence that the remaining workers are more likely to be higher skilled and higher paid if a plant undergoes a major technological upgrade (Bartel, Ichniowski & Shaw 2007, Fernandez 2001).

[8]The literature is unclear whether robots reduce aggregate industrial employment or increase it due to higher demand and a process of highly robotised, more productive firms expanding at the cost of laggards (Acemoglu et al. 2020, Bonfiglioli, Crinò, Gancia & Papadakis 2021). At least the the cross-country level, the former appears more likely, but this is not the central question in this chapter.

[9]The data show that both industrial robots and high-routine workers are concentrated in the manufacturing sector, although there are of course many high-routine workers in sectors that are not heavily affected by robotisation. Nevertheless, it is important to keep in mind that my hypotheses refer to high-routine manufacturing occupations, rather than all routine occupations.

[10]For the purposes of this chapter, it does not matter whether it is indeed the same worker who takes over the more skilled routine job after a technological upgrade or if a better qualified replacement is hired, as long as the 2-digit occupational classification remains the same.

The Fernandez (2001) study investigates detailed changes at one food processing plant. Its findings enrich the picture produced by the sector- and country-wide studies and illustrate some of the mechanisms through which aggregate effects are realised. The study shows increasing skill requirements across the range of production jobs at the plant and in most cases, wage gains relative to the local labour market. At the same time, while the total number of production jobs was essentially constant, the share of highly skilled production jobs (such as maintenance mechanics and electricians) expanded as more complex machinery was adopted. Importantly, these more skilled and better paid production jobs are still in the high-routine occupational groups, raising the average wage of the group as a whole.

The Fernandez (2001) study manages to unpack the black box of technological change and shows how the process plays out on the factory floor. It illustrates the idea of productivity-enhancing upskilling in the wake of robotisation which this study argues may contribute to the increasing relative wage of routine occupations in polarising countries. Furthermore, this effect may be compounded by a negative effect on wages in the low-skill service sector, caused by the influx of former routine workers into service occupations (Dauth et al. 2019). This would further increase the wage gap between high-routine and medium-routine occupations.[11] This leads to the first hypothesis:

H 1: *Higher manufacturing robot density is associated with higher relative wages in routine manufacturing occupations.*

Finding empirical support for this productivity mechanism would be important for several reasons: it would represent a silver lining for workers in these declining occupations, showing that conditional on keeping their job, they stand to benefit financially from greater robotisation. From a policy perspective, finding that robotisation upgrades routine manufacturing jobs would provide a rationale for incentivising robot adoption and thereby the creation of good manufacturing jobs.

[11] On the other hand, Mazzolari & Ragusa (2013) and Moretti (2012) find increasing demand for personal services from high-skill workers which could contribute to higher employment and wages in low-skill jobs.

At the same time, an increase in wage inequality between routine manufacturing and medium-routine service occupations may strengthen the rationale for other interventions such as minimum wage increases.

The second set of hypotheses relates to the impact of unions and employment protection on the relative wage of routine manufacturing occupations.[12] According to power resource theory, labour market institutions increase the bargaining power of workers, and hence their wages (Korpi 1983). Yet, labour market institutions do not affect all workers and occupations equally (Eichhorst & Marx 2015). This seems obvious, but the implications of it have not been studied comprehensively.

The logic of the power resource framework implies that, if in country A 40 percent of workers in occupation O are union members, but only 20 percent of workers in occupation P, workers in O are expected to use their greater bargaining power to negotiate more favourable wage increases for themselves and improve their position relative to workers in P. Similarly, if 60 percent of workers in occupation O in country B are unionised, all else equal, occupation O is expected to be higher up in the occupational wage hierarchy in country B than in country A. Therefore, different occupational groups stand to gain from different institutional setups. Historically, in most countries unions have been stronger in manufacturing occupations than in non-professional service occupations (Card et al. 2004). My argument regarding the effect of union density relies on this still being the case, although recent changes in union membership patterns may invalidate this assumption.[13] If the assumption holds, I expect higher country-level unionisation rates to be beneficial for routine

[12] While minimum wages are another plausible factor influencing relative wages, they are not considered here because many countries in the empirical model do not have a minimum wage. The question of whether minimum wages affect the manufacturing wage premium deserves its own study, especially in light of evidence that high minimum wages increase automation and decrease routine employment (Lordan & Neumark 2018).

[13] Recently, service and public sector occupations in many countries have become relatively more unionised (see, e.g., Bryson et al. 2011, Bureau of Labor Statistics 2021, OECD 2017). In the United States, for instance, unionisation rates in the public sector have overtaken the private sector, although most manufacturing industries also exhibit above-average union membership (Bureau of Labor Statistics 2021). Thus, it is no longer self-evident that high union density should predominantly benefit manufacturing workers, and identification of a differential effect of unions on manufacturing and non-manufacturing wages may not be possible with the available data.

manufacturing wages, but since it cannot be verified with the available data, this hypothesis should be treated as tentative.

Similarly, if half of all workers in occupation O are on temporary contracts and the other half on permanent contracts in both countries A and B, but in country A the temporary workers are protected by an 8 week period of notice whereas there is no such protection for temporary workers in country B, then temporary workers in country A can use the greater fixed cost of firing to extract concessions from their employer, narrowing the wage gap with permanent workers and thus increasing the average wage of occupation O relative to country B. As non-standard employment is most common in low- to medium-skilled service occupations, stricter employment protection for temporary contracts should reduce the manufacturing wage premium, as the occupations that benefit are concentrated in the denominator of the variable. Conversely, since workers in manufacturing occupations overwhelmingly have permanent contracts, we might expect a positive effect on the manufacturing wage premium, but because non-standard employment is more concentrated in specific occupations than permanent employment, I expect the effect of EPL for temporary contracts to be more pronounced. Hence the second set of hypotheses is:

H 2.1: *Higher union density is associated with higher relative wages in routine manufacturing occupations.*

H 2.2: *Stricter employment protection for temporary workers is associated with a lower manufacturing wage premium.*

Moreover, a growing number of voices speculate that labour market institutions may cushion or alter the impact of robotisation on employment (see, e.g., Chiacchio et al. 2018, Dauth et al. 2019, Mishel & Bivens 2017). If there exists statistical moderation with regard to employment, it appears likely that a similar relationship would prevail with regard to wages. I therefore explicitly investigate a potential moderating effect of LMI on the impact of robotisation on wages. The literature suggests that workers will use their bargaining power to influence decisions over new technology in a way that protects their interests (see, e.g., Fernandez 2001).

Thus, in the presence of strong unions, robots might be associated with a higher manufacturing wage premium, as unions negotiate retraining of workers to improve their productivity and protect employment (Fernandez 2001). Conversely, in the presence of strict employment protection for temporary workers, robotisation should boost the manufacturing wage premium by less than it otherwise would, as EPL increases the relative wages of the groups in the denominator. With regard to possible interaction effects, my hypotheses are therefore:

H 3.1: *The (positive) association between robot density and the manufacturing wage premium is stronger in country-years with stronger unions.*

H 3.2: *The (positive) association between robot density and the manufacturing wage premium is weaker in country-years with stronger employment protection for temporary workers.*

Figure 7.2 illustrates the key elements of this theoretical framework. It shows the independent variables on the left, the postulated mechanisms in the middle, and the possible outcomes regarding the wage hierarchy on the right.[14] Since I have argued in chapter 6 that there is a link between occupational wage hierarchies and patterns of employment change, this study investigates what determines wage hierarchies in the first place. Both chapters in conjunction therefore contribute to unpacking the black box of technological change by clarifying some of the related processes. They form the empirical core of this thesis which provides the first account of these often-overlooked intricacies of the technological change debate. Figure 7.2 also highlights that there are additional relevant factors that may influence the manufacturing wage premium and which I control for where possible.[15]

[14]The possible moderating relationships are not shown to keep the illustration straightforward.

[15]Robotisation is an independent variable in my model of relative wages, but it is also an element of RBTC which may affect employment in routine occupations. This complicates the interpretation of the overall effect of RBTC on employment change, but does not undermine any conclusions relating to the manufacturing wage premium which is the focus here.

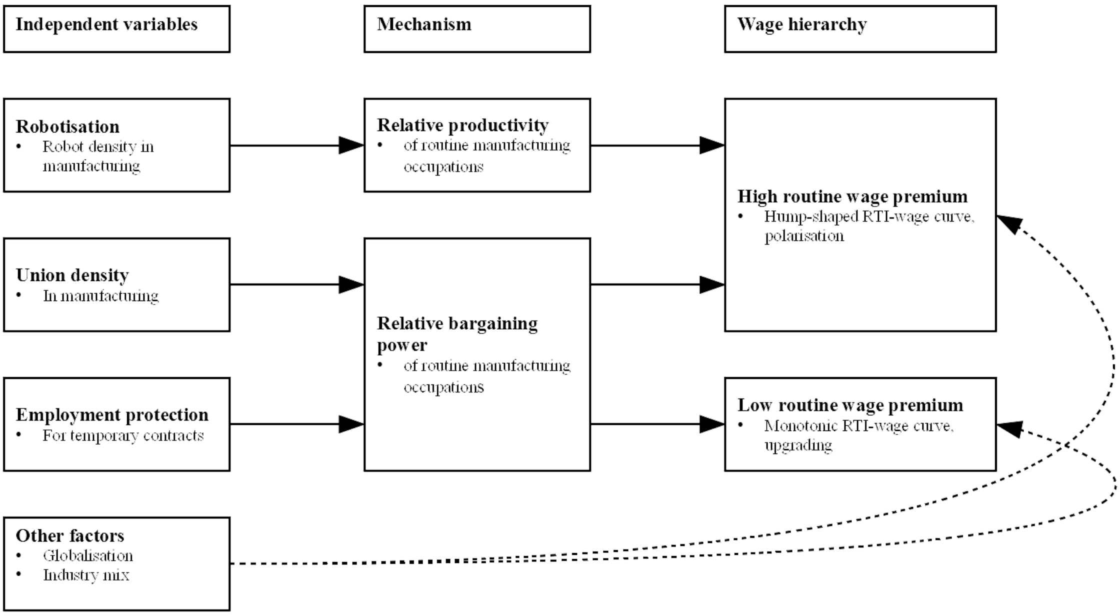

Figure 7.2: *Theoretical framework.*

7.4 Data and methods

7.4.1 Occupational, wage, and institutional data

For the analyses in this study, I have compiled a dataset based on the LIS, IFR, and CPDS datasets, covering the period from 1993 until 2016. The period of analysis is extended compared to chapter 6 to fully utilise the IFR dataset; some further changes to the sample are necessary due to data constraints. Luxembourg does not report robot data to the IFR and the institutional variables are not available for Chile, hence they cannot be included in the main analyses. Simply excluding these two countries would result in a sample of 76 country-years which is too small for meaningful analyses. Therefore, I include additional countries for which LIS, IFR, and CPDS data are available, but which are not necessarily included in all relevant LIS waves. This results in a panel of 150 country-years, based on observations from up to 20 OECD countries.[16] The panel is highly unbalanced, but provides sufficient statistical power for my country-level analyses. Excluding one country

[16] See tables D.2 and D.3 in appendix D for an overview of included country-years and correlations between the variables.

at a time does not qualitatively change the estimates. The main advantages and disadvantages of the LIS, as well as my approach to harmonising the occupational data at the 2-digit ISCO-88 level, have been discussed in detail in chapter 4. I include all workers aged 16 – 65 in the analyses for this chapter, and continue to work with the annual personal labour income variable because it is the most consistently available measure of income in the LIS.

Data on labour market institutions and control variables are mainly taken from the Comparative Political Data Set by Armingeon et al. (2020).[17] The key institutional variables are union density and employment protection for temporary contracts. Ideally, to analyse hypothesis 2.1, union density in the manufacturing sector or better still, the density differential between the manufacturing and service sectors, should be used. I perform some analyses using sector-specific union density data from the Data Base on Institutional Characteristics of Trade Unions, Wage Setting, State Intervention and Social Pacts (ICTWSS) by Visser (2015), but even after interpolation the missingness is too high for these analyses to be meaningful. Therefore, I use country-level union density from the CPDS instead. The assumption here is that unions are strongest in manufacturing and that higher overall union density reflects a bigger density differential between manufacturing and non-manufacturing. This assumption may not hold everywhere, and in any case, relying on country-level data for a sector-level hypothesis will lead to less precise estimation. As some country series contain missing values, the CPDS provides an interpolated version of the union density measure which is preferred here to reduce sample attrition.

The other institutional variable of interest is the level of employment protection. I use the index for temporary contracts, as this is what I expect to have the strongest impact on the wages of non-professional service workers and hence the denominator of the dependent variable. Control variables include the share of manufacturing workers, openness to trade and openness to cross-border capital transactions. These variables are taken from the CPDS database, with some missing

[17]Data on union density and employment protection for a number of missing country-years have been added manually from the OECD database.

Table 7.3: Summary statistics

Variable	Obs.	Mean	Std.Dev.	Min	Max
Log wage ratio	150	1.018	0.01	0.988	1.045
Robot density	150	4.23	3.229	0	13.356
Union density	150	27.829	18.315	5.7	80.845
EPL (temporary)	150	1.32	0.949	0.25	4.5
Trade openness	150	95.75	49.848	20.045	215.438
Capital openness	150	0.945	0.159	0.165	1
Manufacturing employment	150	0.255	0.062	0.151	0.42

observations manually updated with data from the OECD. Summary statistics for the 150-country sample are presented in table 7.3.

7.4.2 Robot data

I view robot penetration as a proxy for the technological advancement of the manufacturing sector in an economy. Data on industrial robots are taken from the International Federation of Robotics, covering the period from 1993 - 2016. This dataset is the gold standard when it comes to the analysis of the aggregate impact of robots on jobs and is used for example by Graetz & Michaels (2018), Dauth et al. (2019), and Klenert et al. (2020). To account for differences in country size, I divide the number of industrial robots by the number of people employed in the manufacturing sector.[18] As my theoretical argument pertains specifically to robotisation in manufacturing occupations, this approach is preferrable to using the number of robots in all sectors and total civilian employment. Following Klenert et al. (2020), I calculate robot densities for the regression analyses using sectoral employment in a base year (1993 or earliest year available) as the denominator to avoid endogeneity issues.

[18]The IFR data on robot stocks assume a robot service life of 12 years. Some authors have instead preferred to use an annual depreciation rate, usually of 10 percent (Graetz & Michaels 2018, Klenert et al. 2020). However, I see no reason to deviate from the approach used by the expert body on industrial robots as well as several other researchers (Acemoglu & Restrepo 2020*a*, Chiacchio et al. 2018).

Despite being the undisputed gold standard for cross-national data on industrial robots, the IFR data pose one important challenge for comparative analyses. In some countries, the IFR only starts to classify robots by sector after 1993. This is for example the case in the US where the sector is not specified at all until 2004 and the percentage of robots in undeclared sectors declines steadily from 89 to 10 percent between 2004 and 2016. Many countries follow a similar pattern. In such cases, using only the reported manufacturing robots would understate the level of robotisation in the earlier period while overstating the growth in robotisation in the later period, and thus bias the measure. This problem affects approximately half of all countries in my sample to varying degrees.

This is recognised by Graetz & Michaels (2018) who therefore interpolate industry-level robot deliveries by multiplying the number of robots reported as "unspecified" by the average share of an industry's deliveries in total deliveries during the years when the breakdown was reported in the data. Chiacchio et al. (2018) likewise state that they allocate unclassified robots based on the distribution of the classified units. Klenert et al. (2020) do not discuss this issue, suggesting that they may be using biased robot numbers. Thus, while most comparative studies account for the failure of the IFR data to adjust for unallocated robots, this requires making some *ad hoc* assumptions which are rarely discussed in the literature.[19] It is clear from inspecting the data, however, that some adjustment for the number of unclassified robots is required in comparative studies.

In this study, I first calculate for each country-year the share s_{cy} of manufacturing robots m_{cy} of the total of the classified robots r_{cy} in that country-year: $s_{cy} = \frac{m_{cy}}{r_{cy}}$. I then calculate a country-specific adjustment factor R_c which takes the average of s_{cy} over the number of years $Y = 1, ..., n$ for which s_{cy} is available:

$$R_c = \frac{\sum_{y=1}^{n} s_{c1} + ... + s_{cn}}{Y}. \tag{7.2}$$

[19] Graetz & Michaels (2018) explain their procedure in the online appendix to their article; Chiacchio et al. (2018) provide no information about their approach.

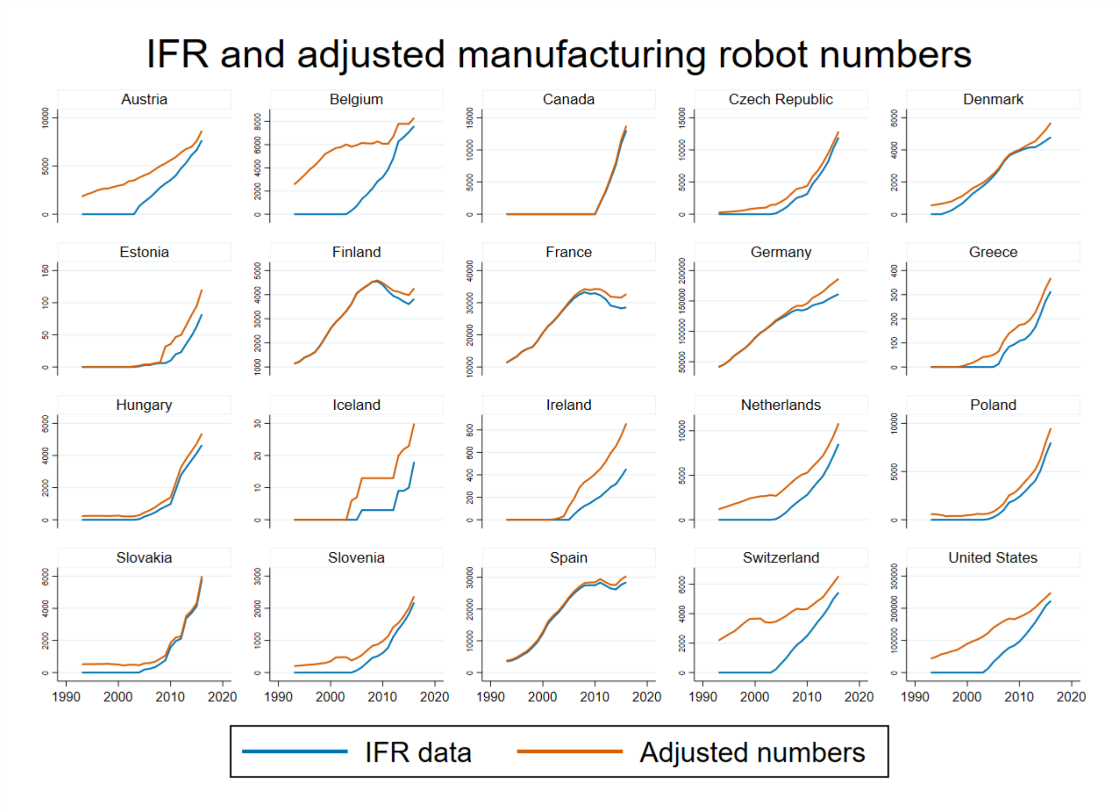

Figure 7.3: *Comparison of IFR and adjusted manufacturing robot numbers by country.*

Then, if u_{cy} is the number of unclassified robots in country c in year y as reported by the IFR, the corrected number of manufacturing robots M_{cy} will be

$$M_{cy} = m_{cy} + R_c u_{cy}. \tag{7.3}$$

This approach is justified since s_{cy} in countries with low shares of unclassified robots is generally quite consistent over time and resembles those in countries where the sectoral classification only starts later. Therefore, it can be assumed that the distribution of classified robots approximates the distribution of unclassified robots. This procedure is very similar to that of Graetz & Michaels (2018), with the difference that I use robot stocks instead of deliveries and apply the procedure to all countries in the sample.[20] Implementing this approach, figure 7.3 shows that

[20]Since Graetz & Michaels (2018) assume an annual depreciation rate of 10 percent instead of following the IFR who assume a fixed life span of 12 years for industrial robots, they calculate their robot stocks from the IFR data on deliveries. Since I follow the IFR's methodology on robot life spans, robot stocks are the appropriate metric to use in my calculations.

in individual countries M_{cy} (orange line) can at times deviate significantly from m_{cy} (blue line). As I will explain in the analyses that follow, performing this adjustment influences the results in the regression analyses; in particular, empirical support for hypothesis 1 vanishes. Yet, the existing literature makes clear that scholars engaged in comparative research have no choice but to perform this or a similar adjustment.

7.4.3 Analytical strategy

I first provide descriptive results that show simple bivariate relationships and patterns, followed by a more thorough econometric analysis. The basic model is an OLS model with country and wave fixed-effects, as the unbalanced nature of the panel leads to problematic levels of attrition in time-series cross-section analyses and makes such models unreliable.[21] The OLS approach yields the same coefficients and standard errors as a fixed-effects panel model in a balanced panel setting, but is more flexible in a real-world unbalanced panel. The model for hypotheses 1 to 2.2 therefore looks as follows:

$$P_{cy} = M_{cy}\beta_1 + U_{cy}\beta_2 + E_{cy}\beta_3 + Z_{cy}\hat{\beta}_4 + \alpha_c + \alpha_y + \varepsilon_{cy}. \tag{7.4}$$

P_{cy} denotes the manufacturing wage premium, M_{cy} is the measure of robot density, and U_{cy} and E_{cy} are measures of union density and employment protection, respectively. The vector Z_{cy} contains further country-level control variables (share of manual workers, trade as a percentage of GDP, capital openness) while α_c and α_y are country and survey wave dummies. I apply robust standard errors clustered at the country-level.

A positive and significant coefficient β_1 would suggest that higher robot density is associated with a higher relative wage of routine manufacturing occupations compared to medium- to high-routine non-manufacturing occupations, supporting hypothesis 1. A positive and significant coefficient β_2 would support hypothesis 2.1 that higher rates of union membership raise the relative average wages of routine

[21] This is compounded by the fact that the LIS data are annual for some countries but not for others. Using only one dataset per wave throws out many degrees of freedom from countries with annualised time-series, whereas using all years would result in the loss of all non-annualised series.

manufacturing occupations. A significant negative coefficient β_3 would provide evidence in favour of hypothesis 2.2 which posits that stricter EPL for temporary contracts reduces the wage premium for routine occupations.

To investigate hypotheses 3.1 and 3.2, an interaction term $M_{cy}{*}U_{cy}\beta$ or $M_{cy}{*}E_{cy}\beta$ is added to the equation, which for hypothesis 3.1 takes the form:

$$P_{cy} = M_{cy}\beta_1 + U_{cy}\beta_2 + E_{cy}\beta_3 + M_{cy} * U_{cy}\beta_4 + Z_{cy}\hat{\beta}_5 + \alpha_c + \alpha_y + \varepsilon_{cy}. \tag{7.5}$$

Hypothesis 3.2 is analysed analogously. Thus, a significant coefficient β_4 in this model would signal that LMI moderate the association between robot density and the manufacturing wage premium as postulated in hypotheses 3.1 and 3.2. Given this nature of the models, the findings will of course not allow for a causal interpretation, but they will point out possible mechanisms that may impact occupational wages and inequality. Statements regarding the "effect" of the independent variables should be interpreted in this vein.

Compared to other articles on the impact of robotisation, this study contributes a more comprehensive geographical coverage and a new perspective on important labour market trends. Most other papers only cover one country and often a shorter time period (e.g. Acemoglu et al. 2020, Dauth et al. 2019). Papers that consider a range of countries such as Graetz & Michaels (2018) do not consider the role of labour market institutions and potential interactions between LMI and robotisation. Moreover, the approach pursued here sheds light on the often-neglected issue of occupational wage hierarchies. Thus, this chapter closes a gap left by the existing body of research in economics and economic sociology.

7.5 Occupational wage hierarchies and country characteristics: Some descriptive evidence

7.5.1 Relative wages of high-routine occupations

First, I describe the manufacturing wage premium which captures the relative wages of high-routine manufacturing occupations compared to medium- and high-routine

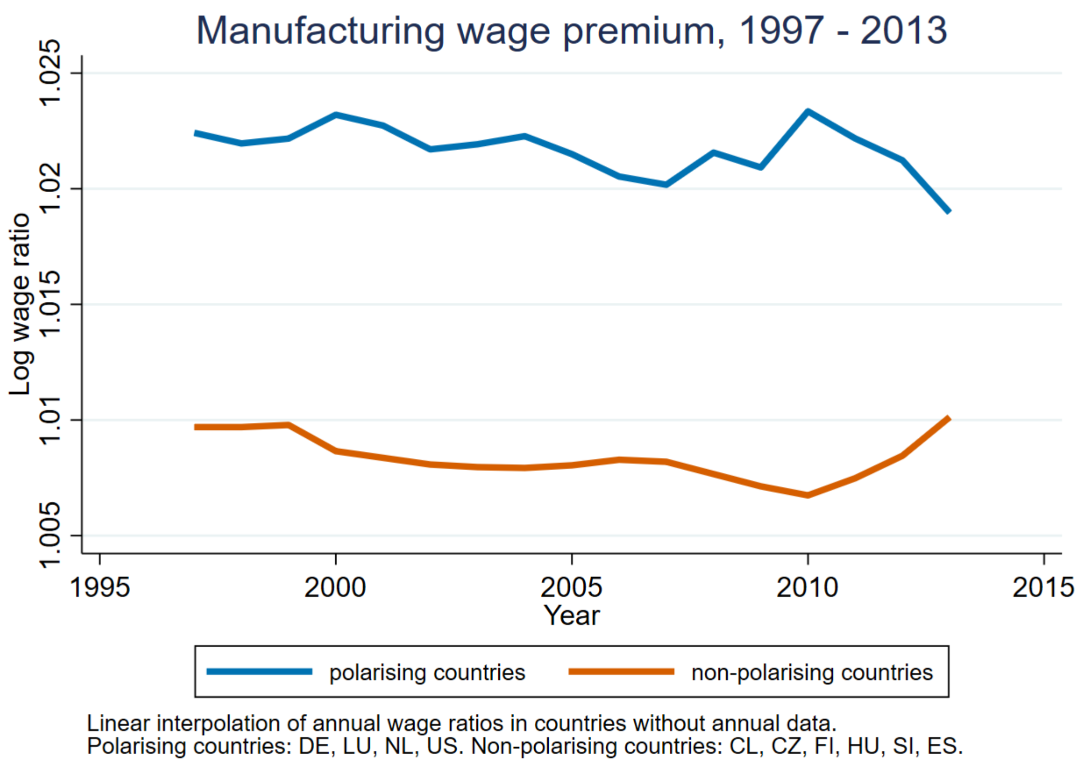

Figure 7.4: *Evolution of the manufacturing wage premia in countries with hump-shaped and monotonic routine-wage curves. The manufacturing wage premium is consistently higher in the hump-shaped (polarising) countries. This validates the measure, as it captures numerically the key difference between the different types of routine-wage curves.*

non-manufacturing occupations. The first important point to note is that there is indeed a manufacturing wage premium in the 20-country sample: the average log wage ratio across countries and years is 1.018 (see table 7.3), corresponding to an average wage premium of 19 percent in routine manufacturing occupations over the comparison group. Moreover, the idea behind this wage ratio is to express the finding from chapter 6, that routine-wage curves come in two distinct patterns, in a simple number. This means that the manufacturing wage premium should be higher in countries with a hump-shaped RTI-wage curve, as the hump-shaped pattern is the graphical manifestation of higher relative wages for high-routine workers.

Figure 7.4 shows that indeed the ratio is consistently higher in the 4 countries with hump-shaped curves compared to the 6 countries with monotonic curves from the earlier analyses, thereby validating my measure of relative wages.[22] The values correspond to a 25 percent manufacturing wage premium in the hump-shaped countries and only a 6 percent premium in the monotonic countries in 2010. Since a hump-shaped RTI-wage curve is predictive of employment polarisation, this also implies that in polarising countries high-routine manufacturing workers earn higher relative wages. Overall, the descriptive evidence shows the usefulness of my measure of relative wages, as it corresponds well to the routine-wage curves and hence the patterns of employment polarisation and upgrading highlighted in chapter 6. It is therefore suitable to investigate the research questions of this chapter, which centre around the question why occupational hierarchies differ between countries in the first place.

7.5.2 Robots and routine wages

Overall, the IFR data show that robot density has increased substantially everywhere since the mid-1990s. There are also, however, wide (and widening) disparities between countries regarding robot penetration, as shown in figure 7.5. The figure shows the evolution of manufacturing robot density in 20 OECD countries. While in Germany in 2016 the number of industrial robots per 1000 industrial workers is 16.5, the runners up Denmark are far behind with about 10.8 robots per 1000 workers and in most countries the number is substantially lower still.

In figure 7.6, we see that between 1993 and 2016, in the countries where routine occupations cluster near the middle of the wage distribution, robot penetration in the manufacturing sector has increased at a markedly faster pace than in countries where routine occupations earn lower wages. Furthermore, robot adoption picked

[22]This analysis is based on the 10 countries with an uninterrupted LIS time-series from wave 4 to 9 that were analysed in chapter 6. There, Germany, Luxembourg, the Netherlands, and the United States showed a hump-shaped RTI-wage curve and polarising employment change. Meanwhile, Chile, the Czech Republic, Finland, Hungary, Slovenia, and Spain all exhibited the monotonic RTI-wage curve and, except in Hungary, upgrading employment change. This relationship holds if additional countries with long LIS time-series, such as Belgium and Ireland, are included. Note that in the shorter timeframe of this chapter, the employment trend in the United States looks more like polarised upgrading while in Finland employment change appears to be downgrading, compared to the earlier pattern in chapter 6 of straightforward polarisation and upgrading, respectively. This can be seen in figure D.1 in appendix D.

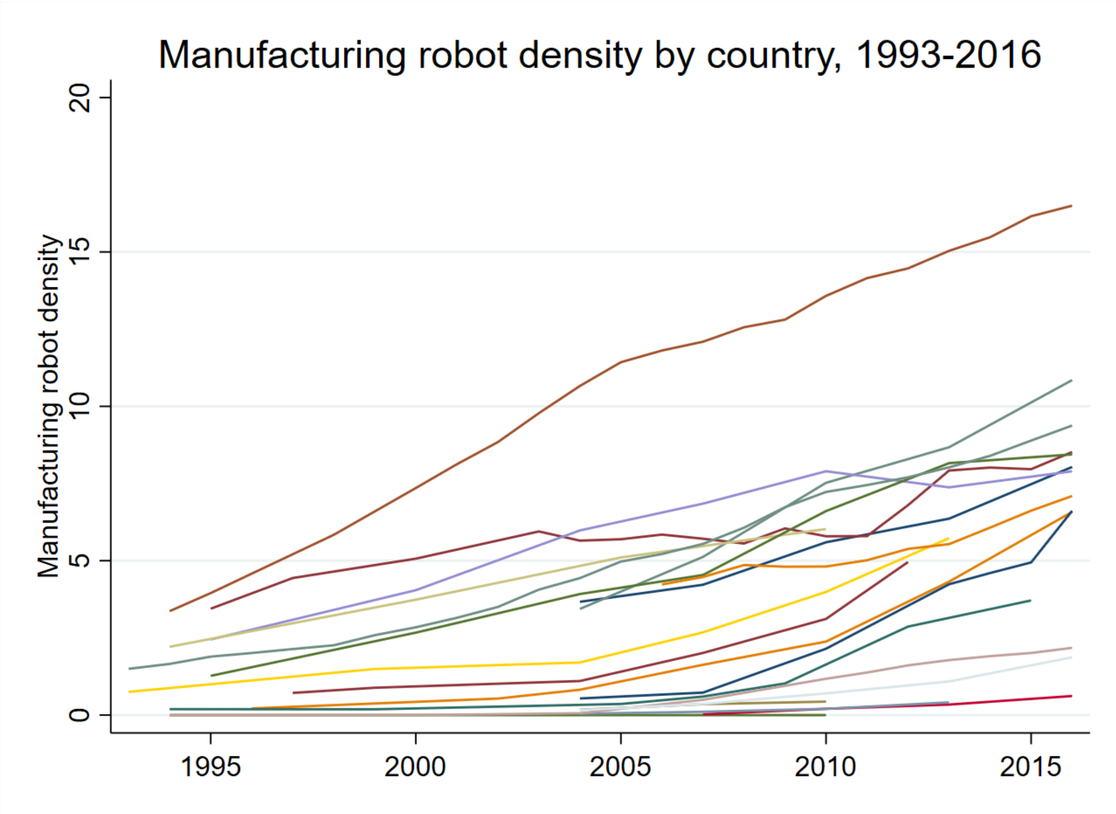

Figure 7.5: *Manufacturing robot density by country, 1993 - 2016.*

up in the countries where the monotonic routine-wage curve gave way to the hump-shaped one while continuing to lag behind in those with a consistently monotonic curve. Thus, the countries where robot adoption has been fastest are also where routine workers earn higher relative wages. This is a hint that robotisation may improve the relative wage position of routine occupations in more advanced countries and thereby contribute to employment polarisation there as posited in hypothesis 1.

Is greater robot density really associated with higher relative wages in routine manufacturing occupations which are susceptible to robotisation? Unsurprisingly, robot density is positively correlated with the absolute real wages of all occupation groups, presumably reflecting the role of economic development which drives both variables. However, robot density is also positively correlated with the manufacturing wage premium in 5 out of 7 LIS waves, with a stronger positive correlation in later waves. This indicates that not only are absolute real wages higher in countries with higher robot density, but that robotisation is specifically associated with higher

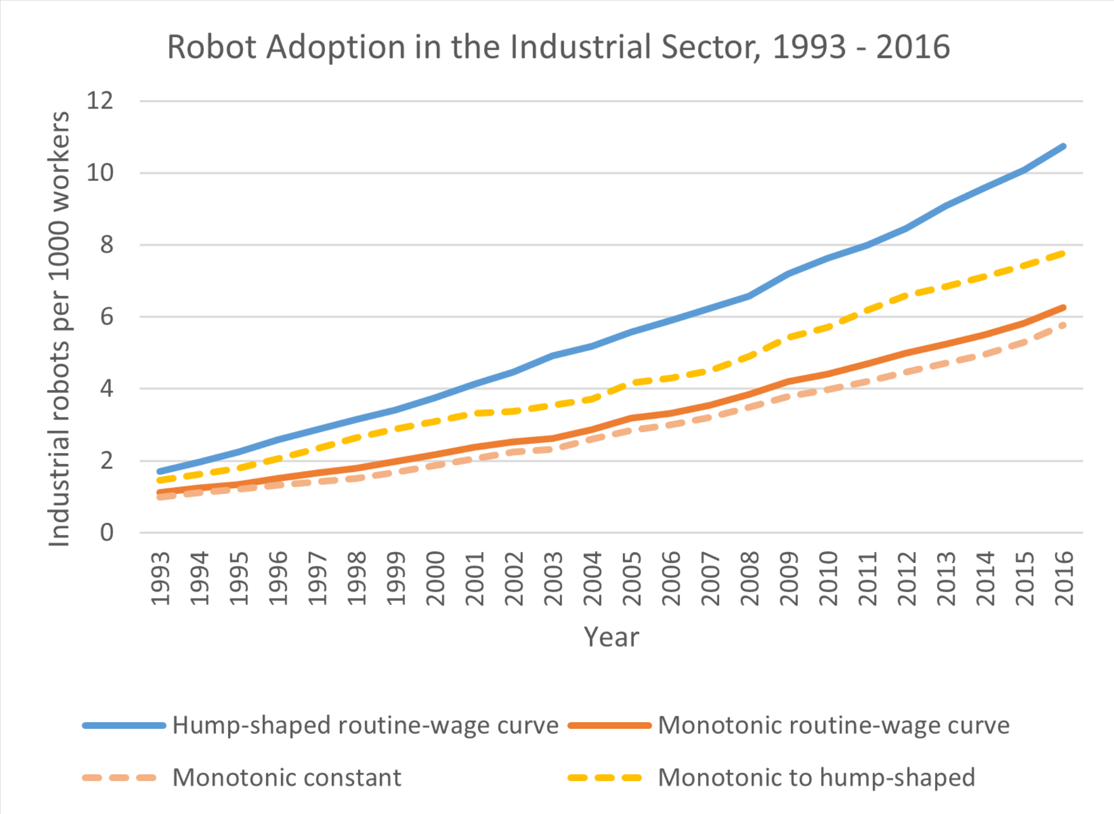

Figure 7.6: *Robot penetration in different RTI-wage clusters; all countries with robot and sectoral employment data, 1993 – 2016. Hump-shaped: BE, DE, DK, NL, US. Monotonic: AT, CH, CZ, ES, FI, FR, HU, PL. Monotonic to hump-shaped: ES, CH. Monotonic constant: AT, CZ, FI, FR, HU, PL.*

relative wages for high-routine manufacturing workers.

Figure 7.7 plots the correlation for the earliest and latest available waves; considering the intermediate waves (see figure D.2 in the appendix), there appears to be a clear trend towards a more positive relationship. However, the correlations are not statistically significant, partly due to the low N, and should therefore not be overinterpreted. Moreover, this analysis does not control for any confounding factors. Thus, the correlation may reflect a common underlying cause and, if there is a causal relationship, its direction is unclear. Hence, more thorough research is needed. Despite the larger number of cases (N=93), the relationship between changes in robot density and changes in the manufacturing wage premium is not statistically significant and is, in fact, weakly negative, as can be seen in figure D.5. Overall, the descriptive analyses regarding industrial robots so far largely

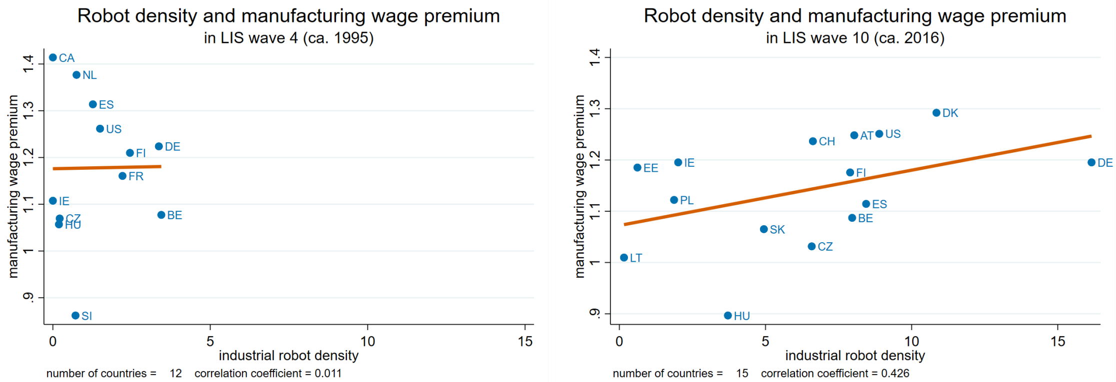

Figure 7.7: *Correlation of robot density and manufacturing wage premium, 1995 and 2016.*

point in the direction of the effect postulated in hypothesis 1. Manufacturing robot density has increased, most strongly in countries which exhibit a high manufacturing wage premium, and the correlation between robot density and the manufacturing wage premium is positive and has increased. Only a slight negative relationship between changes in robot density and changes in the wage premium indicates that the hypothesised productivity mechanism may not operate as expected.

7.5.3 Labour market institutions and wages

The second set of hypotheses in this study is that unionisation and employment protection influence wage hierarchies by boosting the wages of some occupations relative to others. Unions are expected to boost the relative wages of routine manufacturing workers, while EPL for temporary workers should raise the average wages in low-wage non-manufacturing occupations. This section explores some descriptive evidence regarding these hypotheses. 17 countries in my sample provide information on unionisation or EPL in both 1993 and 2016. Of these, 7 countries each exhibit a hump-shaped routine-wage curve/a high manufacturing wage premium and the remainder a monotonic routine-wage curve/a low manufacturing wage premium.

As we see in table 7.4 and figure 7.8, union density has declined markedly. In the mid-1990s the average unionisation rate was very similar in both groups of countries at around 45 percent. In the following two decades, the data reflect the widespread decline in union membership that has been covered extensively elsewhere (see, e.g., Bryson et al. 2011). Interestingly, however, the decline has been much more pronounced in the countries with a monotonic routine-wage curve. The countries with hump-shaped curves where manufacturing wages have fared better also now have a more unionised workforce than the countries with monotonic curves which have predominantly faced occupational upgrading. This is compatible with the union wage effect found by Parolin (2021) and my hypothesis 2.1, that unionisation should benefit the relative wages of routine manufacturing workers. The boxplot for 2016 furthermore shows some convergence in the monotonic countries, the outlier being Finland with a unionisation rate of 65 percent.[23] Overall, there is a pattern of uneven union decline and convergence which is largely consistent with hypothesis 2.1.

With regard to employment protection, the small overall decline between 1993 and 2016 hides marked differences between the hump-shaped/polarising and monotonic/upgrading countries. In the latter, the EPL index for temporary contracts has essentially been constant, although there, too, has been some convergence and the countries with the strictest legislation have liberalised their regulations somewhat. The median EPL score in these countries in 2016, however, is higher than in 1993. In the hump-shaped/polarising countries, on the other hand, employment protection for non-standard workers has been slashed. The average score of 1.02 in 2016 is 37 percent lower than the 1.61 in 1993, and the median level of EPL has declined as well. Moreover, there has been a great deal of convergence, so that the 2016 values for the 7 hump-shaped countries fall in a fairly narrow range between 0.25 and 2.06.[24] These patterns are well in line with the theoretical expectations formulated in hypothesis 2.2.

Taken together, the descriptive findings for this sample of 17 mainly European countries are only partly in line with Gallie (2017, p. 238) who argues that there was "a stronger tendency towards linear skill upgrading in [EU] countries with more

[23]Across all waves and countries, figure D.3 in the appendix shows that while the relationship between union density and the manufacturing wage premium is negative in the 1990s, a positive relationship prevails in the 2010s. This is largely due to the inclusion of a cluster of post-Soviet countries which combine low unionisation and a low manufacturing wage premium. The findings in table 7.4 and figure 7.8 are not affected by this change in sample composition, however.

[24]Across all waves and countries, the relationship between employment protection and the manufacturing wage premium is consistently negative, as shown in figure D.4 in the appendix.

Table 7.4: Unionisation and EPL scores by country cluster, 1993 and 2016

	Unionisation		**Employment protection**	
RTI-wage curve	1993	2016	1993	2016
Hump-shaped (N=7)	42.6	31.6	1.61	1.02
	(20.6)	(20.9)	(1.67)	(0.69)
Monotonic (N=10)	46.3	18.0	1.86	1.84
	(26.0)	(17.5)	(1.48)	(0.61)
Overall (N=17)	44.8	23.6	1.75	1.66
	(23.3)	(19.6)	(1.51)	(0.76)

Note: Standard deviations in parentheses. The sample includes all countries for which the CPDS provides data in both years. The samples for union density and EPL are almost identical: Canada and Greece are included in the latter instead of Denmark and Estonia in the former, all other countries are identical.

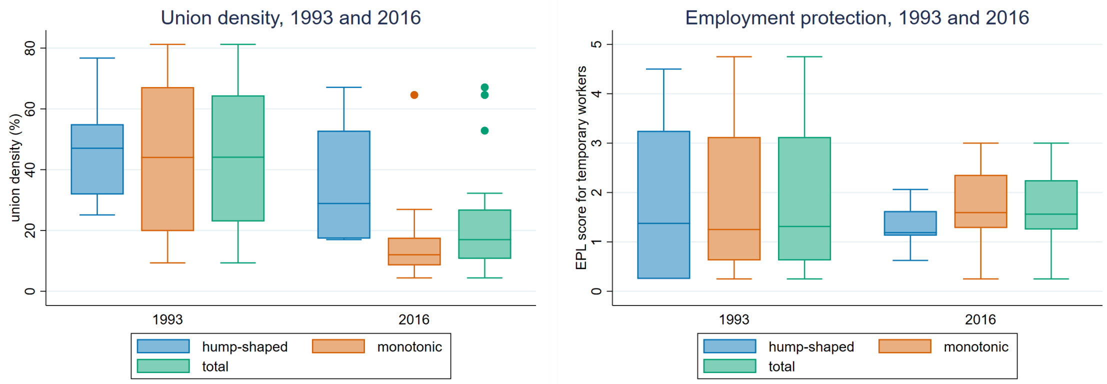

Figure 7.8: *Changes in union density and employment protection in hump-shaped/polarising and monotonic countries, 1995 and 2016.*

egalitarian institutions". At least regarding unions, my theory and analyses so far point in a different direction, as unions have been more stable in polarising countries. This has to do with my inclusion of the Eastern European countries, of which many have seen a rapid decline in union density in the 1990s, as figure 1.8 showed. However, employment protection for temporary workers appears to be associated with a lower manufacturing wage premium as both my theory and Gallie (2017) predict.

In general, the descriptive analyses have produced mixed evidence. The manufacturing wage premium has been validated as a measure of relative wages and robot density appears to be associated with a higher manufacturing wage premium. With regard to labour market institutions, the descriptive evidence is clear for employment protection but more ambiguous for union density. While every effort has been made to keep the samples as consistent as possible, the comparability of these descriptive analyses is limited as a consequence of the various sources that are employed for data on wages, robots, and institutions. Thus, multivariate analyses on a common sample are needed before stronger conclusions can be drawn.

7.6 Regression analyses

7.6.1 Main model

To further investigate the impact of robot density and labour market institutions on occupational wage ratios, I now turn towards the formal estimates of the hypotheses. Table 7.5 presents results for a basic model including the focal variables as well as country and wave dummies.

Model 1 shows just the three focal variables, robot density, union density, and EPL. Hypothesis 1 is not supported, as a higher number of industrial robots per 1000 manufacturing workers in 1993 is not associated with a better relative wage position of high-routine manufacturing workers. Likewise, the coefficient on union density is marginally negative and not statistically significant, indicating that high levels of unionisation may not in fact boost the relative earnings of routine manufacturing occupations. Hypothesis 2.1 is therefore not supported by this simple model. Only hypothesis 2.2, that stricter employment protection for non-standard workers reduces the wage premium for high-routine manufacturing occupations, receives empirical support.

The absence of any statistically significant association of robot density and union density is surprising. In the case of union density, this may have to do with the coarse country-level data which tell us nothing about the sectors in which union members are concentrated. The insignificant but negative estimate of robot density indicates that the hypothesised productivity effect may be held in check by a countervailing displacement effect. The significant negative coefficient on EPL supports my reasoning that, since temporary contracts are much more common in the medium-routine occupations in the denominator of the dependent variable, stricter regulation of such contracts should raise the wages of those occupations relative to high-routine manufacturing occupations where non-standard work is less common.

In the remaining models, I include interaction terms between robot density and union density and robot density and EPL to investigate hypothesis 3.1, which states that robot density should be more positively related to the manufacturing wage premium where union density is higher, and 3.2, according to which stricter EPL should mute the expected positive relationship between robot density and relative manufacturing wages. In both cases, the analyses in columns 2 and 3 offer at best weak support for the hypotheses. Although the coefficients on the interaction terms point in the expected directions, neither the main effects nor the interaction terms are anywhere close to statistical significance. In column 4, the interaction terms are statistically significant at the 10 percent level and moreover, the main effect for union density turns positive. All in all, however, the evidence from this simple model is too weak to conclude that LMI influence the relationship between robotisation and the relative wages of routine manufacturing occupations.

In column 1 of table 7.6, I estimate the main model based on equation 7.4 which includes my set of control variables. An overview of the country-years included in the analysis can be found in table D.2 in appendix D.[25] Trade openness and capital openness are included to account for the arguments that link globalisation with polarisation through various mechanisms such as outsourcing, "defensive innovation", or increases in overall inequalities (Acemoglu 2003, Burgoon 2001, Michaels et al. 2014, Thoening & Verdier 2002). Trade openness, defined as the sum of imports and exports as a percentage of GDP, has no association whatsoever with the manufacturing wage premium. However, we see a highly significant negative

[25]It is worth noting that due to data limitations the panel is heavily unbalanced, with Germany and the US accounting for 19 and 24 of the 150 total observations in the main model, respectively. However, as remarked above, including only one observation per wave would unduly reduce the degrees of freedom. Moreover, I have verified that excluding Germany or the US from the analysis does not affect the conclusions.

Table 7.5: Determinants of the manufacturing wage premium, no controls

DV: log manufacturing wage premium	(1)	(2)	(3)	(4)
Robot density (RD)	-0.0705	-0.0970	-0.0073	-0.0220
	(0.061)	(0.060)	(0.066)	(0.058)
Union density (UD)	-0.0002	-0.0015	0.0036	0.0033
	(0.024)	(0.024)	(0.023)	(0.021)
EPL	-0.2837**	-0.2613**	-0.1210	-0.0203
	(0.112)	(0.123)	(0.175)	(0.170)
RD * UD		0.0021		0.0032*
		(0.001)		(0.002)
RD * EPL			-0.0500	-0.0704*
			(0.039)	(0.038)
Observations	165	165	165	165
R-squared	0.821	0.825	0.826	0.834

Robust standard errors (clustered at country level) in parentheses. All models include a constant and country and wave dummies. The dependent variable has been multiplied by 100 for better readability of the coefficients. *** $p<0.01$, ** $p<0.05$, * $p<0.1$

coefficient on capital openness, suggesting a lower premium on high-RTI employment in countries that are more open to cross-border capital transactions.

The model also controls for industry employment as a share of total civilian employment, which essentially captures the importance of the manufacturing sector. The coefficient is positive and statistically significant, showing that where manufacturing is economically important, relative wages in the sector tend to be higher. This undermines the narrative that robotisation leads to fewer but better paid routine manual workers: to the contrary, where there are relatively more industry workers, their position in the wage hierarchy is more favourable.[26] [27]

[26] In light of this, manufacturing employment would be another plausible candidate for an interaction effect with robot density, but additional analyses produce no evidence of such a relationship (results available upon request). Thus, it appears that higher industrial employment is associated with a higher manufacturing wage premium independent of the degree of robotisation of an economy.

[27] I performed additional analyses with controls for the average RTI of the workforce in a country-year, real GDP growth, and the share of part-time employment. None of these variables showed any relationship with the dependent variable. In the case of part-time employment, it is important to note that this dispels concerns that the use of annual earnings and differences in

Table 7.6: Determinants of the manufacturing wage premium, with controls

DV: log manufacturing wage premium	(1)	(2)	(3)	(4)
Robot density (RD)	-0.1276**	-0.1323**	-0.0632	-0.0698
	(0.060)	(0.059)	(0.057)	(0.054)
Union density (UD)	-0.0492*	-0.0495*	-0.0394	-0.0387
	(0.028)	(0.028)	(0.027)	(0.027)
EPL	-0.2625**	-0.2590**	-0.1056	-0.0654
	(0.118)	(0.120)	(0.141)	(0.150)
RD * UD		0.0004		0.0014
		(0.002)		(0.002)
RD * EPL			-0.0503**	-0.0589**
			(0.019)	(0.023)
Trade openness	0.0023	0.0022	0.0030	0.0027
	(0.005)	(0.005)	(0.005)	(0.000)
Capital openness	-3.4408***	-3.4295***	-3.1730***	-3.0846***
	(0.955)	(0.957)	(0.944)	(0.966)
Industry employment	8.6123**	8.5462**	9.2806**	9.1476**
	(3.472)	(3.497)	(3.750)	(3.839)
Observations	150	150	150	150
R-squared	0.874	0.874	0.878	0.880

Robust standard errors (clustered at country level) in parentheses. All models include a constant and country and wave dummies. The dependent variable has been multiplied by 100 for better readability of the coefficients. *** p<0.01, ** p<0.05, * p<0.1

Furthermore, after controlling for trade openness and capital openness, the coefficients on robot density and union density become statistically significant, although only at the 10 percent level in the case of unionisation. Comparing the models in table 7.5 and table 7.6, a Ramsey RESET test for omitted variables strongly suggests the presence of omitted variable bias in the former, but not in the latter. The negative coefficient on both explanatory variables indicates a lower premium on routine work in country-years with high robot density or a highly unionised workforce. These findings conflict with my hypotheses 1 and 2.1 which

the prevalence of part-time work might distort the wage ratios. Results from these analyses are likewise available upon request.

each predicted a positive association.

With regard to unions, this most likely reflects the changing patterns of union membership and the insufficiently detailed resolution of the data. While traditionally, the industrial sector has experienced the strongest wage effects from unionisation (Card et al. 2004, DiNardo et al. 1996), membership patterns have changed and public sector and highly educated workers may increasingly reap the benefits of unionisation (Bureau of Labor Statistics 2021, OECD 2017). Given that single-country studies with more granular data such as Parolin (2021) tend to find strong relationships between union membership and wages, my finding should not be taken to question the importance of unions as such, but rather to illustrate the importance of context and data quality.

The compressing effect of robots is more difficult to explain. Graetz & Michaels (2018) found robot density to be positively linked to the wages of medium-skilled and manufacturing workers, and according to Dauth et al. (2019), for manufacturing workers who stay with the same employer, robotisation is associated with a substantial wage premium. By contrast, in my sample, the evidence points towards the opposite relationship between robot density and the relative wages of routine manufacturing workers. Of course, the different studies are not directly comparable because they look at different subgroups and metrics. Importantly, with the LIS employment data I am restricted to analyses at the country-level rather than the country-sector-level. Nevertheless, the findings of studies such as Graetz & Michaels (2018) informed the hypotheses in this chapter, which I am however unable to substantiate. It appears that across countries and over time, the displacement effect of robotisation, according to which a decrease in an occupation's employment share should also depress its wages, outweighs the wage gains from productivity improvements (Acemoglu & Autor 2011).

It is worth pointing out again that the coefficients on robot density and union density are only statistically significant in the models which also control for capital openness. Furthermore, the size of the coefficients varies considerably between the models in tables 7.5 and 7.6. Because of the direction of the coefficients, it is clear that hypotheses 1 and 2.1 are not supported; rather, there is some evidence for an opposite effect. Only the coefficient on EPL is consistently statistically significant and stable across model specifications, in line with hypothesis 2.2. In terms of variance explained, the partial η^2 for model 1 in table 7.6 shows that robot density and union density each explain approximately 7.4 percent of the total variance, whereas EPL accounts for approximately 5.5 percent. All the main independent variables are dwarfed, however, by the partial η^2 of 21.4 percent for capital openness.

The size of the estimated effects is nevertheless quite substantial, as can be seen in figure 7.9. A standard deviation increase in robot density is associated with a decrease in the manufacturing wage premium by 43 percent of one standard deviation. Union density shows a stronger picture: a one standard deviation increase is linked to a decrease in the manufacturing wage premium by 93 percent of a standard deviation. This comparison creates the impression that union density might be the quantitatively more impactful factor, but it is important to remember that union density within a country is less variable than robot density, which has been rapidly expanding in many countries.[28] For example, in Germany between 1994 and 2016, union density has fallen by less than one standard deviation from 30 percent to 17 percent, while robot density has increased by 3.24 standard deviations from 3.37 to 13.84. A standard deviation increase in employment protection for temporary workers is associated with a decrease in the manufacturing wage premium by 26 percent of a standard deviation. Even more so than union density, employment protection is quite stable over time and in fact does not change at all in many countries, with most variation being between countries.

The analyses in columns 2 - 4 of table 7.6 estimate the interaction model from equation 7.5. Column 2 shows no evidence whatever that union density moderates the relationship between robot density and the manufacturing wage premium, as all coefficients are virtually unchanged compared to column 1. However, unlike in the model without controls, in column 3, I do find a significant moderating effect of EPL on the relationship between robot density and the manufacturing wage premium.

[28] This is to say, a larger share of the total variation in union density is between countries, whereas with robot density, there is a lot of within-country variation.

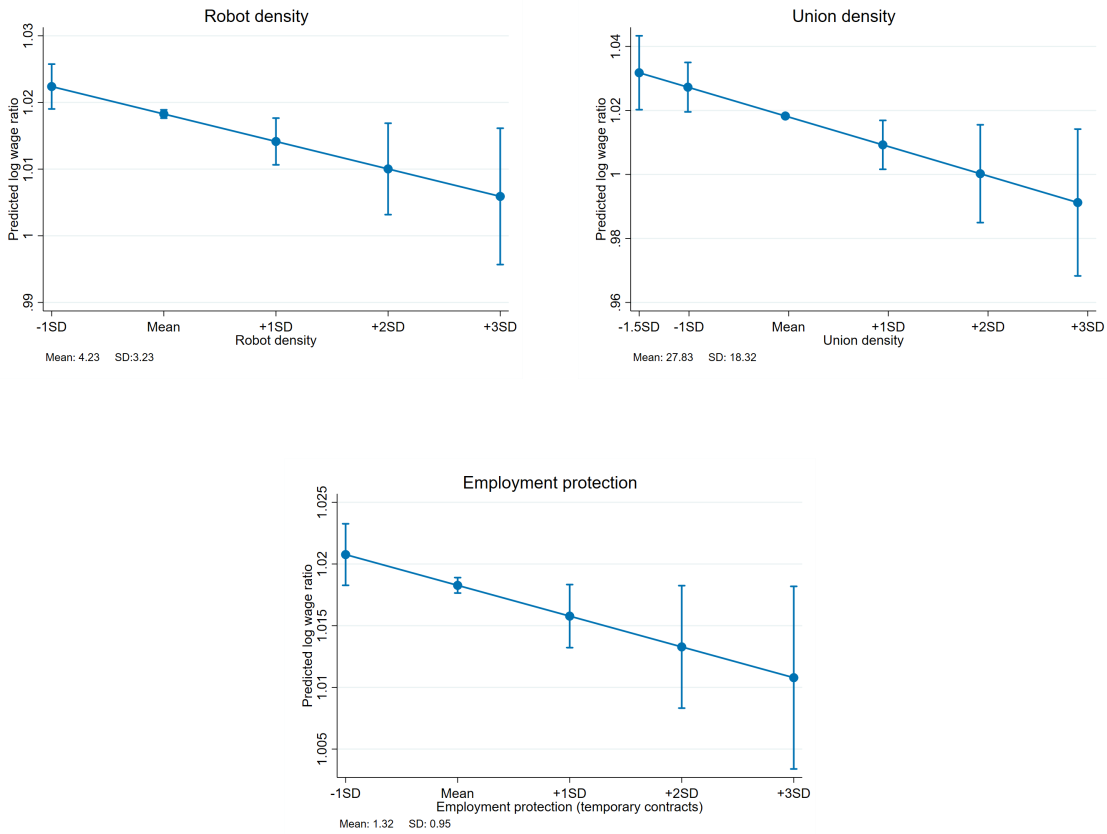

Figure 7.9: Marginal effects of the key explanatory variables based on the model in column 1 of table 7.6.

This provides partial support for my hypothesis 3.2. While the effect of robot density is always negative, it is more strongly so where employment protection is stricter. The same result holds when both interaction terms are included jointly in column 4.

To visualise how the effect of robot density is moderated by the level of employment protection for temporary contracts, figure 7.10 plots the predicted values of the manufacturing wage premium for different values of robot density and EPL, based on the model in column 4. The figure shows that in the region where most observations are located, with EPL scores between 1 and 3 and robot densities between 1 and 8, this interaction is indeed substantively meaningful. An equivalent figure for the insignificant interaction between union density and robot density is included in appendix D. Taken together, my main model finds no support for my hypotheses relating to robot density and union density. However, employment

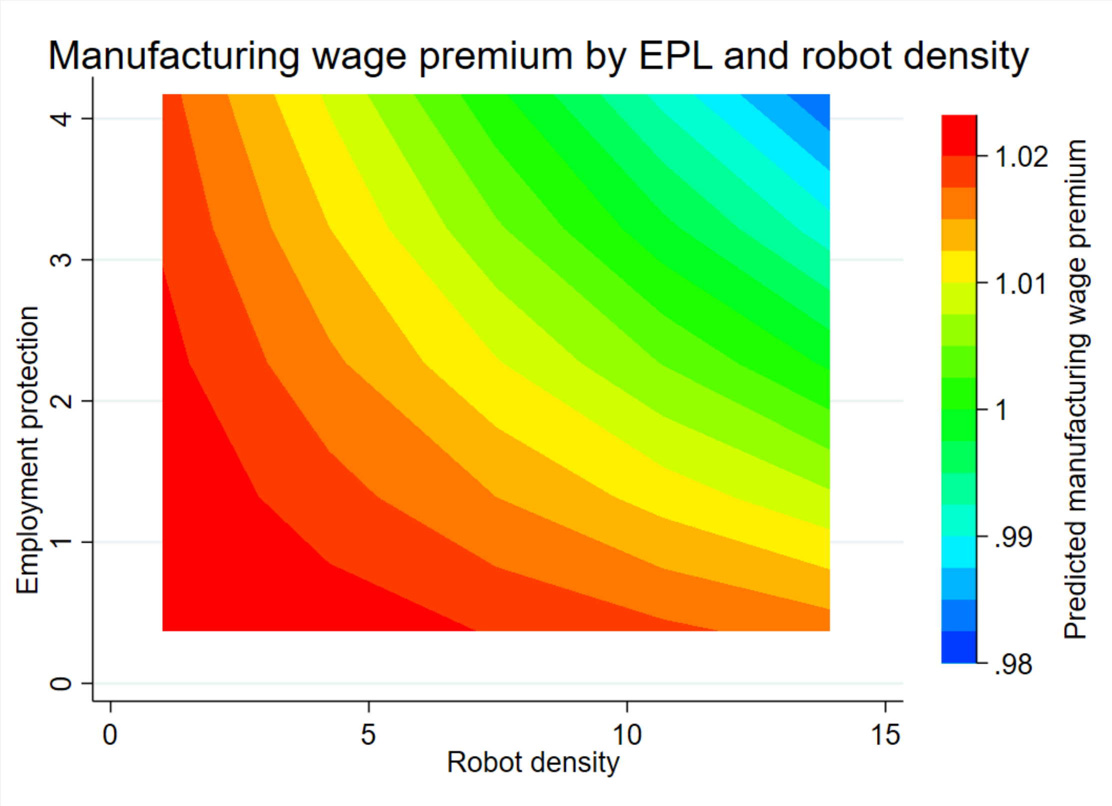

Figure 7.10: *Predicted manufacturing wage premium by EPL and robot density based on the model in column 4 of table 7.6.*

protection for temporary contracts is shown to exert an independent effect on the manufacturing wage premium (hypothesis 2.2) as well as to moderate the (otherwise insignificant) effect of robotisation (hypothesis 3.2).

7.6.2 Additional analyses

This section reports additional analyses that deal with some of the limitations of this study. Table 7.7 investigates relationships in different subsamples. In column 1, I account for the nature of the routine-wage relationship by interacting the independent variables with a dummy for the hump-shaped RTI-wage curve. Eight countries in the regression sample are classified as belonging to the hump-shaped/polarising group, while the remaining twelve have a predominantly monotonic routine-wage curve associated with occupational upgrading.[29] In the model,

[29]The two country clusters are: hump-shaped: Belgium, Canada, Denmark, Germany, Iceland, Ireland, Netherlands, United States; monotonic: Austria, Czech Republic, Estonia, Finland, France, Greece, Hungary, Poland, Slovakia, Slovenia, Spain, Switzerland.

the dummy itself is positive but not statistically significant, while the interaction terms are negative in all cases and significant for union density and EPL. This suggests that higher union density and stricter EPL have a stronger dampening effect on the manufacturing wage premium in countries with a hump-shaped routine-wage curve than in countries with a monotonic one.

This is also borne out in the analyses in columns 2 and 3, where the hump-shaped and monotonic samples are analysed separately. It appears that the significant coefficients in the combined model are largely the result of significant relationships in the hump-shaped sample. It may be that the distinction between time-series and cross-sectional variation comes into play here. The hump-shaped cluster contains 89 observations while the monotonic cluster only accounts for 61. This results in an average number of observations per country in the hump-shaped cluster of 11.1, compared to only 5.1 observations per country in the monotonic cluster. As a result, the estimation in the hump-shaped sample addresses more time-series variation, while the variation in the monotonic sample is to a greater extent cross-sectional. Once more LIS data series become annualised, a larger and more balanced panel will eventually allow for more sophisticated analyses using methods such as error correction models (Beck 2001, De Boef & Keele 2008).

In column 4, only one observation per country and wave is included. This reduces the overweight of countries like Germany and the US in the model and results in the loss of over one third of all observations from the main model. The results point in the same direction as in the main model, but the negative coefficients on robot density and EPL are not statistically significant, while union density, capital openness, and industry employment retain very similarly sized and significant (in case of union density, at the 10 percent level) coefficients. To further alleviate concerns that the findings may be driven by individual countries, I have verified that the results of the main model do not change if one country is excluded at a time. Thus, the analyses of different subsamples in table 7.7 reveal some interesting patterns, in particular with regard to differences between hump-shaped/polarising and monotonic/upgrading countries.

Table 7.7: Analyses by sub-sample

	Hump-shaped interaction	Hump-shaped only	Monotonic only	1 obs/wave only
DV: log manufacturing wage premium	(1)	(2)	(3)	(4)
Robot density (RD)	-0.1401*	-0.0734**	-0.1720	-0.0750
	(0.068)	(0.029)	(0.167)	(0.071)
Union density (UD)	-0.0280	-0.0663**	-0.0530	-0.0570*
	(0.025)	(0.020)	(0.034)	(0.028)
EPL	0.1759	-0.2136*	0.2996	-0.1527
	(0.210)	(0.095)	(0.351)	(0.159)
Hump-shaped = 1	1.9778			
	(1.159)			
Hump-shaped * RD	0.0177			
	(0.060)			
Hump-shaped * UD	-0.0623**			
	(0.023)			
Hump-shaped * EPL	-0.5621**			
	-0.253			
Trade openness	-0.0027	0.0054	-0.0112	-0.0007
	(0.006)	(0.004)	(0.010)	(0.007)
Capital openness	-2.5046***	0.3384	-3.0277**	-3.3093***
	(0.806)	(0.953)	(1.105)	(0.907)
Industry employment	8.4410*	12.4884**	5.6457	7.5347**
	(4.504)	(4.294)	(7.936)	(3.453)
Observations	150	89	61	98
R-squared	0.885	0.862	0.872	0.885

Robust standard errors (clustered at country level) in parentheses. All models include a constant and country and wave dummies. The dependent variable has been multiplied by 100 for better readability of the coefficients. *** p<0.01, ** p<0.05, * p<0.1

Table 7.8: Analyses using alternative independent variables

	Manufacturing UD	Bargaining coverage UD	Regular EPL	Unadjusted RD
DV: log manufacturing wage premium	(1)	(2)	(3)	(4)
Robot density	-0.2472**	-0.0556	-0.0907	0.1392**
	(0.104)	(0.084)	(0.056)	(0.062)
Union density	-0.0009	0.0035	-0.0485	-0.0386
	(0.028)	(0.009)	(0.029)	(0.031)
EPL	0.5115	-0.2271	1.1397***	-0.2110
	(0.557)	(0.132)	(0.364)	(0.180)
Trade openness	0.0059	0.0059	0.0038	-0.0017
	(0.008)	(0.008)	(0.005)	(0.006)
Capital openness	-4.3537**	-1.4871	-3.6089***	-4.2449**
	(1.719)	(0.881)	(0.871)	(1.873)
Industry employment	10.3757	8.0872	6.8506	5.4313
	(6.950)	(5.460)	(4.752)	(3.872)
Observations	71	119	150	116
R-squared	0.893	0.907	0.876	0.892

Robust standard errors (clustered at country level) in parentheses. All models include a constant and country and wave dummies. The dependent variable has been multiplied by 100 for better readability of the coefficients. *** p<0.01, ** p<0.05, * p<0.1

In table 7.8, I re-estimate the main model using different operationalisations of the independent variables. This illustrates the importance of aligning theory and measurement, similar to chapter 5. One concern that I have referred to repeatedly is the quality of the union density data. Using a country-wide measure of union density to investigate a hypothesis that is really about the role of unions in the manufacturing sector is clearly suboptimal. Ideally, union density in the manufacturing sector or better still, the density differential between the manufacturing and service sectors, should be used. Unfortunately, comprehensive sectoral data are not currently available. The ICTWSS by Visser (2015) is the most comprehensive source for union-related data. It reports union density by sector, but the coverage is very sporadic: only for 23 country-years in my main sample, manufacturing union density is available. By interpolating missing data, it is possible to increase the total number of country-years to 71, but with less than one third of the data points being actual observations, the findings of this model in column 1 must be treated with caution. Unsurprisingly, the results are less reliably estimated; furthermore, EPL now shows a positive coefficient. This signals that no reliable conclusions can be drawn based on the sectoral union density data that are available. It remains therefore unclear whether hypothesis 2.1 has been rejected because insufficiently detailed data were used to test it or because there genuinely is a negative association between union density and the manufacturing wage premium.

In column 2, the adjusted bargaining coverage variable of Visser (2015) is used instead of union density. This variable captures the proportion of all wage and salary earners with the right to bargain who are covered by collective bargaining agreements. The estimate is positive but far from statistical significance, as are the coefficients on the other independent variables. It is also an order of magnitude smaller than the estimate for union density in the main model, strengthening the case for the importance of unions as such. Column 3 shows that EPL for regular contracts has a highly significant positive relationship with the manufacturing wage premium, in accordance with the bargaining power argument that EPL for standard contracts should above all raise wages in occupations with very low rates of non-standard work, such as routine manufacturing occupations. This, of course, is the flip side of my hypothesis 2.2. I formulated hypothesis 2.2 as I did because I expected a stronger impact of EPL for temporary contracts based on the reasoning that the benefits should be more concentrated in specific occupations. This analysis suggests, however, that the power resource mechanism is also clearly visible with regular EPL.

Finally, in column 4, I use the unadjusted robot data (m_{cy} instead of M_{cy}) to illustrate that it is crucial to correct the IFR robot data. Without the adjustment, the coefficient on robot density is positive, statistically significant, and sizeable.[30] Thus, failure to correct for the number of unallocated robots could have led us to erroneously fail to reject hypothesis 1. However, this would have been based on

[30]The lower number of observations is due to country-years with zero manufacturing robots in the original data, such as the US before 2004, being dropped from the analysis.

clearly inaccurate assumptions such as that there were zero manufacturing robots in the United States prior to 2004. Comparative studies that appear to overlook this limitation, such as Klenert et al. (2020), may reach conclusions that do not withstand closer scrutiny. Conversely, studies which do correct for unallocated robots should clarify their procedure and how this affects their findings. All in all, the analyses in table 7.8 show that the results in this chapter are sensitive to different operationalisations of the independent variables. At the same time, however, they validate my theoretical arguments for why the measures used in the main analyses are most appropriate for investigating my hypotheses.

7.6.3 Limitations

There are several remaining limitations to this study. First and foremost, the manufacturing wage premium is necessarily an imperfect measure of a complex phenomenon such as occupational wage hierarchies. Any attempt to express the key facets of a distribution in a single number entails substantial trade-offs. In the present study, I focus on high-routine occupations in the manufacturing sector. This allows me to investigate not only the role of labour market institutions but also of robot density, which affects the manufacturing sector first and foremost. However, other high-routine occupations which in this approach are part of the denominator may also be influenced by RBTC more broadly. Thus, it is important to be clear that the goal of this chapter has not been to quantify the impact of overall technological change on wage hierarchies, but of robotisation, alongside and in relation to institutional factors.

There are some further limitations related to the data used in the study. The necessary adjustment of the robot numbers that has already been discussed at length, as has the nature of the union data. Besides this, I see two main areas of concern. For one, the mismatch between occupational routine-intensity scores and industry robot densities is likely to muddy the findings. With wage and routine-intensity data at the country-sector-occupation level, it would be possible to allocate workers more accurately to the high-routine manufacturing and the medium-to-high-routine non-manufacturing groups. Once the required data for such an analysis become available, a reduction of this type of measurement error should lead to more precise estimates and allow for more confident conclusions.

Furthermore, like in chapter 6, the use of annual labour income means that the values of the dependent variable may be affected by differences between countries and occupations in the prevalence of non-standard work. In additional analyses, I found no evidence that national differences in the prevalence of part-time work drive the results. However, the share of part-time work at the national level is again an imperfect measure of a phenomenon where much of the variation is at the sectoral or occupational level.

Finally, the highly unbalanced panel precludes the use of time-series cross-section methods to better account for the time component in the model (see, e.g., Beck & Katz 1995, De Boef & Keele 2008, Keele & Kelly 2006). As more annualised series become available in the LIS, more sophisticated analyses will become possible in due course. Keeping these limitations in mind, I nevertheless believe that this study provides an important new perspective on the question of what determines occupational hierarchies and patterns of employment change.

7.7 Discussion

The aim of this chapter was to determine whether, to the extent that higher relative wages for routine manufacturing workers are associated with employment polarisation, heterogeneous effects of robotisation and LMI on the wages of routine manufacturing and low- to medium-skill service workers contribute to these wage differences. To this end, I investigated the independent and moderating effects of robotisation, unionisation, and EPL on the wage premium for routine manufacturing occupations.

This represents the logical continuation of the previous research in this thesis. Contrary to earlier arguments (see, e.g., Goos et al. 2014), the decline of employment in high-routine occupations has not led to pervasive employment polarisation across developed countries, but to a mix of polarising and upgrading employment change

(Fernández-Macías & Hurley 2017). My research in chapter 6 has found that the position of high-routine occupations in the wage hierarchy is an important predictor of the pattern of occupational change which a country experiences. This chapter has shown that robotisation and labour market institutions, both independently and jointly, influence the wage premium for routine manufacturing occupations compared to other medium- to high-routine occupations. This makes it possible to better explain the patterns found in chapter 6.

The study investigated two sets of mechanisms and hypotheses. The first mechanism relates to the impact of the adoption of industrial robots on the relative productivity of routine manufacturing workers. Based on previous research in labour economics and sociology on the effects of robotisation (see, e.g., Fernandez 2001, Graetz & Michaels 2018), I hypothesised that robots increase the relative productivity of the routine workers who work alongside them, and hence the their relative wages.

The second mechanism is rooted in the welfare state and comparative institutionalist literatures which emphasise the importance of worker bargaining power (Korpi 1983, Korpi & Palme 1998, Visser & Checchi 2011). Where institutions place more power in the hands of certain workers, they should be able to negotiate more favourable wages. I therefore postulated that higher union density (stricter employment protection for temporary contracts) should strengthen (weaken) the relative bargaining power of routine manufacturing workers vis-à-vis medium- to high-routine non-manufacturing workers, with the accompanying implications for wages.

Furthermore, I hypothesised that unionisation and employment protection moderate the relationship between robotisation and wages. I expected a stronger positive relationship where unions are strong, and a weaker (albeit positive) relationship in country-years with stricter EPL. I investigated these hypotheses using data from the LIS, IFR, and CPDS for a panel of 150 country-years covering the period from 1993 until 2016.

The productivity hypothesis with regard to robotisation, despite favourable evidence in descriptive analyses, was not supported by the regression analyses.

Manufacturing robot density has increased, most strongly in countries which exhibit a high manufacturing wage premium, and the correlation between robot density and the manufacturing wage premium is positive and has increased. However, in a panel model controlling for institutional variables and other potential confounders the effect of robotisation on the manufacturing wage premium is consistently estimated to be negative. This suggests that in the absence of rapid robotisation in many polarising countries, the routine manufacturing premium would be higher than it is now. In upgrading countries, a slower pace of robotisation could have contributed to polarisation. This also indicates that some other factor may be simultaneously exerting upward pressure on routine wages and promoting robot adoption. General economic development comes to mind, but I could find no evidence that GDP growth or levels affect the manufacturing wage premium.

Regarding union density, the descriptive analyses showed that unionisation declined across the board, but more so in upgrading countries with a monotonic RTI-wage curve. However, in the regression models, higher union density is associated with a lower manufacturing wage premium, although this relationship is not particularly robust. One possible explanation for this is the changing pattern of union membership. While union density has historically been highest in the manufacturing sector (hence the original hypothesis), the composition of union membership has been shifting towards white collar and public sector workers (OECD 2017, Visser & Checchi 2011). Manufacturing workers may therefore no longer be the primary beneficiaries of high unionisation. A closely related point is that the country-level union data used in the analyses are not strictly speaking suitable to analyse hypothesis 2.1 which pertains to occupational or at least sectoral union densities. Hence, more detailed data would need to be available for a sufficient number of countries to investigate the hypothesis more conclusively.

Employment protection for temporary workers, which has declined in countries with a hump-shaped RTI-wage curve and stayed constant in countries with a monotonic one, is robustly associated with a lower manufacturing wage premium, as postulated by hypothesis 2.2. As non-standard work is rare in the group of high-routine manufacturing occupations in the numerator of the dependent variable, this implies that EPL raises the wage floor for temporary workers who are concentrated in the occupations in the denominator of the dependent variable. The decline of EPL in polarising countries may have counteracted other equalising trends that would have eroded the manufacturing wage premium.

I also investigated whether the impact of robotisation on the manufacturing wage premium is contingent on the level of unionisation and EPL. I found no evidence that higher unionisation is associated with a more positive relationship between robot density and the manufacturing wage premium, leaving hypothesis 3.1 unsupported. However, there is qualified support for hypothesis 3.2, that the (positive) association between robot density and the manufacturing wage premium is weaker in country-years with stronger employment protection for temporary workers. While the relationship is not positive to begin with, the negative association is indeed more pronounced in country-years with strict EPL. Table 7.9 summarises the findings regarding the five hypotheses that I tested in this chapter.

Table 7.9: Summary of the findings

Hypothesis	Finding
H 1	Rejected
H 2.1	Rejected
H 2.2	Supported
H 3.1	Rejected
H 3.2	Partially supported

Overall, only two of my five directional hypotheses receive some empirical support. Robot density and union density are not associated with a higher manufacturing wage premium, and there is no interaction between the two variables. Nevertheless, it would be premature to completely discard the underlying mechanisms. Some of my analyses find statistically significant effects in the opposite of the expected direction. While leading me to reject the respective hypotheses, this can still be interpreted to indicate the presence of the productivity and bargaining power mechanisms - but with different implications than the existing literature and my hypotheses suggest.

This chapter has addressed the final research question of my thesis. It has found some evidence that robotisation and LMI may have heterogeneous associations with the wages of different occupational groups, albeit not always in line with my hypotheses. Therefore, they may contribute to the emergence - or not - of employment polarisation that has been documented in chapter 6. However, further research with better data will be necessary to substantiate both the significant and the null findings. Overall, the chapters of this thesis have shown that with an innovative approach that borrows from both the economics and sociological literatures, important new insights are possible. Yet, there is still much to be learned about the interplay of technology and institutions in the labour market. In the concluding chapter, I synthesise what this thesis has discovered and how its research agenda can be pushed further.

8
Conclusion

Contents

8.1 Summary of the findings

It is my hope that the findings of this thesis will contribute to a better understanding of the effects of technological change and labour market institutions on employment, wages, and inequality in developed countries. Previous research has established broad patterns such as declining employment in routine and low-skilled occupations, increasing overall wage inequality, and a polarisation of employment and wages in several countries. Among the numerous explanations which scholars have put forward for these trends, technological change and the decline of inclusive labour market institutions feature prominently in the literature in labour economics and sociology. However, there has been very little dialogue between the two fields, and existing research has struggled to provide convincing explanations for exceptions to the broad patterns mentioned above.

My thesis has integrated insights from labour economics and sociology to address several blind spots in this literature and therefore will be relevant to researchers in both fields. Based on the overall research puzzle, I formulated three interlinked research questions:

1. How can occupational task content be measured in a way that corresponds to the underlying theoretical concepts and that accounts for the relevant dimensions of variation?

2. Can the diverse patterns of employment change in developed countries be reconciled with the key tenet of RBTC, that technological change everywhere substitutes for routine workers?

3. Do robotisation and LMI, independently and jointly, have heterogeneous effects on the wages of different occupational groups?

I will now revisit each of these questions in turn to assess how my theoretical arguments have held up, highlight the contributions of my research, and place the findings in a wider context.

8.1.1 Measurement of and variations in occupational task content

The first research question was born out of a realisation that existing measures lack solid theoretical and empirical foundations. Yet, correctly identifying to which degree different occupations entail routine tasks is essential for empirical analyses of RBTC theory. For comparative analyses, this also implies that variation in occupational task content, either between countries or over time, deserves greater attention. To address this research question, I developed improved indices of occupational task content which exhibit meaningful differences with established measures as well as across countries and over time.

I used data from the EWCS, in which more than 107,000 respondents from 27 EU countries between 2000 and 2015 answered questions about the tasks they perform in their job, to construct measures of routine intensity and task complexity. The analyses in chapter 5 showed that these measures can lead to important improvements in the operationalisation of occupational task content in the five problem areas that I identified. These improvements include an index of task complexity that can be used to capture SBTC, a better alignment of concepts and measurement, use of survey data to account for within-occupation variation, and the incorporation of geographical and temporal variation.

Most important for this thesis was the better alignment of concepts and measurement. I argued that the items used by Autor et al. (2003) and Autor & Dorn (2013), and to a lesser extent by Fernández-Macías & Hurley (2017), do not properly measure the repetitiveness and codifiability of occupational tasks, which is how routine intensity is commonly defined, including in the papers in question. My comparative analyses in chapter 5 showed that the correlation with the standard measure of routine intensity in the literature, the index of Autor & Dorn (2013), is only 0.74 - a less than impressive correlation considering that they are meant to be measuring the same underlying concept. The crucial difference is the classification of predominantly manual occupations such as machine operators, trades, and elementary occupations. These occupations are generally considered more routine intensive by my index, while clerks, who were considered highly routine-intensive in Autor et al. (2003), receive a lower RTI score. The differences with the indices of Fernández-Macías & Hurley (2017) are found to be much less pronounced, which was expected given that the same data source was used. Nevertheless, the analyses show that my routine and complexity measures constitute further improvements.

The second crucial contribution of chapter 5 was to show the significant degree of variation in task content that exists between countries and over time. This means that the practice of using task data from the US that in some cases date back to the late 1970s, is difficult to defend if more appropriate measures are available. At the very least, researchers should pool data to reduce reliance on a single country. I also pointed out the new opportunities for research based on country- and wave-specific task measures, although due to the composition of my samples, I went on to use the pooled version of the indices in chapters 6 and 7. Overall, chapter 5 showed that a way of operationalising occupational task content that better aligns theory and measurement and accounts for relevant variation is available, thus answering research question 1 and setting up the remainder of the thesis.

8.1.2 How to reconcile upgrading employment change and RBTC?

The second research question was, whether and how the diverse patterns of employment change in developed countries can be reconciled with the key tenet of RBTC theory, that technological change substitutes for routine workers. For if routine workers were indeed medium wage, as posited for example by Goos et al. (2014), their declining employment shares should have ushered in pervasive employment polarisation. However, numerous studies such as Fernández-Macías (2012), Oesch (2013), and Oesch & Piccitto (2019) show this not to be the case.

Critics of the labour economics literature who have pointed out this inconsistency have as a consequence often called into question the routine-biased nature of technological change; my argument in chapter 6, however, took a different route. I argued that in some countries the most routine-intensive occupations are predominantly low-wage, implying occupational upgrading as a result of RBTC. In making this argument, I combined insights from sociology and labour economics in a novel way that allowed me to account for a phenomenon which both disciplines separately have struggled to explain convincingly.

To investigate this argument, I used a 10-country sample of the LIS covering the years 1995 - 2013. The sample included eight European countries, including three Eastern European ones, as well as Chile and the United States, thus reflecting the diversity in terms of size, economic development, and institutional setup of the OECD. For each country-year, I produced a routine-wage curve as a measure of the concentration of high-routine occupations in the wage structure.

In line with my argument, I found that in the countries which had experienced polarisation, the routine-wage curve was predominantly hump-shaped, implying a concentration of high-routine occupations around the middle of the wage distribution. Conversely, in the countries which had shown upgrading employment change, the routine-wage curve was monotonically negative, indicating a clustering of routine occupations at the bottom of the wage hierarchy. Thus, I demonstrated that RBTC is indeed only associated with employment polarisation in countries where routine occupations are concentrated in the middle of the wage hierarchy. This empirical pattern provides strong support for my refined theory of RBTC. The routine-wage curves therefore emerge as a helpful tool for analysing the impact of technological change on occupational change.

The findings of chapter 6 illustrate the importance of comparative research and help further debunk the "myth" of pervasive employment polarisation (Oesch & Piccitto 2019), while offering a plausible alternative explanation that does not discard the RBTC model altogether. Instead, the chapter augmented the RBTC model with the insight from the sociological literature that employment polarisation has not been pervasive and thereby contributed to a better understanding of the effect of technological change on employment.

8.1.3 How to explain different wage hierarchies?

The findings of chapter 6 raised the follow-up question, what explains the different relative wage positions of routine occupations in polarising and upgrading countries. To investigate this question, it was necessary to consider a broader set of contributing factors. Thus, chapter 7 asked if robotisation and LMI, independently and jointly, have heterogeneous effects on the wages of high-routine manufacturing workers and medium- to high-routine non-manufacturing workers.

This question required the introduction of another new concept, the manufacturing wage premium. I defined this as the ratio of the average log wage of high-routine manufacturing occupations over the average log wage of medium- and high-routine non-manufacturing occupations. This measure captures wage inequality between occupational groups in the lower two thirds of the wage distribution and served as the dependent variable in my analyses. It is generally greater than 1, indicating that there is indeed a wage premium for routine manufacturing occupations. This facet of wage inequality has thus far received very little attention, but as chapter 6 made clear, a higher manufacturing wage premium indicates polarising RBTC.

Keeping with my previous approach, I combined insights from sociology and economics to pursue three arguments. First, I argued based on recent research in labour economics that industrial robots may increase the relative productivity and wages of those workers who remain employed in heavily robotised occupations. Thus, my hypothesis was that higher manufacturing robot density should be associated with a higher manufacturing wage premium.

Secondly, I turned to power resource theory and the sociological and welfare state literature to argue that unionisation and employment protection not only affect the overall dispersion of wages, but also the relative wages of occupations based on the differential exposure of occupations to these institutions. Based on historical membership patterns, I argued that where unionisation is higher, routine manufacturing workers would benefit disproportionately from higher union wages and therefore the wage premium should be higher. Furthermore, I postulated that employment protection for temporary workers should reduce the manufacturing wage premium because temporary workers are concentrated in the low-wage service occupations that make up the denominator of the dependent variable.

Finally, I considered the possibility that LMI moderate the impact of robotisation on relative wages. While higher unionisation should accentuate the expected effect of robotisation, stricter EPL should reduce it.

I again relied on the LIS, albeit with an enlarged sample of up to 20 countries, robot data from the IFR, and institutional data from the CPDS. My hypotheses were only partly supported by the data. Most importantly, robotisation and union density were associated with a reduced manufacturing wage premium, although the relationships were not very robust. EPL for temporary contracts was robustly associated with a lower manufacturing wage premium as expected, and it was also shown to amplify the negative effect of robotisation (rather than dampening the anticipated positive relationship).

While these findings are not causal, I interpreted the lower wage premium in more robotised economies as an indication that the productivity mechanism, at least at the country-level, is trumped by the effect of an oversupply of manufacturing workers. The counterintuitive effect of unionisation could be due to changes in union membership patterns, whereas the finding that in more regulated labour markets, robotisation more strongly reduces the routine wage premium, highlights the importance of accounting for institutional factors in analyses of the impact of robotisation.

Overall, the analyses in chapter 7 showed that robotisation and LMI, both independently and jointly, have heterogeneous effects on occupational groups as evidenced by the manufacturing wage premium. However, they cannot explain why the manufacturing wage premium is higher in polarising countries, as all three variables are associated with a lower wage premium. Therefore, further research is needed to fully explain the underlying reasons for the findings of chapter 6.

8.1.4 The effect of technological change and labour market institutions on employment, wages, and inequality

In terms of the overall research puzzle that motivated this thesis, we can conclude that technological change and labour market institutions affect employment, wages, and inequality in numerous ways. This thesis has highlighted some that have been neglected in prior research. The findings thus contribute substantially to a better understanding of trends in the labour markets of developed economies.

By focusing first on task content, then employment, and finally wages and inequality as outcome variables, I have worked along the chain of reasoning of RBTC theory and at every stage contributed new theoretical impulses and empirical insights. The key to doing so was the concerted effort to bridge gaps between the literatures in labour economics and sociology. By adapting the RBTC framework in light of findings from sociology, and explicitly analysing interactions between the two explanatory frameworks, I hope, the finished thesis is more than the sum of its parts.

The main takeaways from this thesis are as follows. The importance of proper measurement of key theoretical concepts was illustrated in chapter 5 and the subsequent application of the proposed measures throughout this thesis. The chapter furthermore showed that there is meaningful variation in occupational task content across countries and over time. Chapter 6 demonstrated the importance of comparative research and of questioning commonly held assumptions by showing that RBTC does not necessarily lead to employment polarisation and is in fact compatible with upgrading employment change that has continued to characterise many countries. Finally, chapter 7 established that robotisation and LMI both influence the relative wages of occupations that are particularly affected by RBTC. Furthermore, EPL may moderate the impact of robotisation on wages. Therefore, robotisation and LMI may affect how RBTC impacts overall employment trends. All in all, this thesis provides a rich set of new insights and manifold starting points for further research into a topic that will only increase in importance as technological change transforms economies.

8.2 Limitations and avenues for further research

Although I believe that this thesis makes important contributions to the scientific literature, it is also clear that its scope is limited and its analyses, incomplete. Mindful of this fact, in this section I outline these limitations and propose avenues for further research that build upon the novel findings and measures presented here.

8.2.1 Parsimony amid a plethora of competing explanations

As I illustrated in chapters 1 and 2, the phenomena that I sought to explain in this thesis have been linked to a plethora of possible explanations which, besides technology and institutions, include globalisation, growing female labour force participation, immigration, and other demographic changes. Thus, the need to impose some restrictions in the interest of parsimony necessarily constitutes a limitation of my studies. Compared to other research in the field, I have sought to provide a more comprehensive picture, but it is important to recognise that there still remains a trade-off with parsimony.

To strike the best balance in this trade-off, I decided to focus on technological change and labour market institutions because these are most consistently mentioned in the literature, while controlling for other factors where possible. Hence, there are additional macro-trends that affect labour markets which have not been studied in detail here. Most notable among these factors are the potential influence of part-time work and female labour force participation on the routine-wage curves that have been discussed in chapter 6, and the possible impact of immigration and other LMI such as minimum wages on the manufacturing wage premia in chapter 7. Incorporating these factors would be a valuable extension of the research in this thesis, which will however be complicated by limitations of the available data.

8.2.2 Data limitations

Indeed, a recurring issue throughout the thesis was the availability and quality of empirical data. Generally, more comprehensive analyses have been impeded by the poor data availability for large-N comparative studies. For example, in chapter 7, I was able to include some globalisation variables, but comparable information on migration streams was not available. Analyses that cover a wider range of variables may therefore have to be limited to smaller samples or even individual countries where large, high-quality datasets are available.

But also the data that are available have numerous shortcomings that limited the depth of my analyses. As I pointed out in chapter 3.4, I combine data from several different datasets and use occupation- and industry-level data, which do however not allow for the construction of more precise industry-occupation-level measures. Moreover, some of the variables that are used in the empirical analyses represent a second-best measure of the underlying concept.

Combining different data sources which cover slightly different geographical areas and time periods leads to high levels of attrition in the sample. This has been particularly evident in chapter 7, where the analyses incorporated LIS, IFR, and CPDS data. As I mentioned in chapter 7.6, the unbalanced panel precluded a more sophisticated analysis of the temporal variation in the data.

Another important limitation of the data in this thesis is that while I use occupation-level data from the EWCS and LIS and industry-level data from the IFR, I am unable to construct more precise industry-occupation-level measures of task content, employment, wages, and robot density, due to either insufficient sample sizes or missing or incompatible information. While based on the findings of Goos & Manning (2007) and Acemoglu & Autor (2011) I argued in chapter 3.4 that this is unlikely to qualitatively affect my results, reliable industry-occupation measures would have allowed for more precise estimation in chapters 6 and 7.

Finally, several of the variables used in the empirical analyses are a second-best measure of the respective concept. While the task measures developed in chapter 5 offered a superior alternative to the widely used measures of Autor & Dorn (2013), the use of annual labour income and country-level union density was necessitated by a lack of data on hourly wages and sectoral union densities, and thus was itself a second-best choice which entails problems which I discussed at length in chapters 6 and 7. Thus, while this thesis has provided important insights based on a wide range of empirical data, its broad comparative approach has not been without drawbacks in terms of data availability and quality.

8.2.3 Towards a full use of my novel task measures

Another important limitation of this study - which at the same time is an obvious call for further research - is the fact that I did not make full use of the possibilities afforded by my novel task measures. Part of my reason for calling for these new measures was the lack of country-specific or time-varying measures in the wider literature. Yet, in chapters 6 and 7, I limited myself to using the pooled measure that is akin to the existing indices of Autor & Dorn (2013) and Fernández-Macías & Hurley (2017). Hence, the critical reader may ask, why bother with the country- and wave-specific indices in the first place?

First of all, as explicated above, the analysis in chapter 5.7 of changes in task intensity and differences between countries yielded important findings in its own right. However, the presence of non-European countries in the analyses in chapters 6 and 7 would have required using country-specific measures in European countries and pooled measures in non-European countries, which would be confusing and inconsistent. I therefore used the pooled measures throughout. As chapter 5.6 has shown, the pooled measures still have important advantages over Autor & Dorn (2013) and Fernández-Macías & Hurley (2017), such as better capturing the underlying concept of routine-intensity and using worker survey data. Moreover, showing that even my pooled measure, which is constructed analogously to the existing ones, yields new insights, in no way diminishes the importance of pursuing country-specific or time-sensitive analyses where appropriate. The detailed measures will be most fruitfully applied to European datasets such as the EU-SES in analyses that focus on more recent time periods (the EU-SES starts in 2002). I have not been able to do so within the confines of this thesis; however, future research will surely benefit from more closely examining country differences and temporal changes in task content, as captured by my measures.

8.3 Whither technological change - is this time different?

Notwithstanding its limitations, this thesis has produced numerous insights about labour market changes in the recent past. Even though it has produced no evidence for large-scale displacement of workers or dramatic increases in inequality due to technological change, the potential implications of future technological disruption will undoubtedly remain a focal point for scholars. In particular, many researchers worry that "this time is different", and that with advances in machine learning and artificial intelligence, technological unemployment may be imminent (Mokyr et al. 2015). Thus, even though this thesis was not concerned with forecasting the susceptibility of occupations to future technological change, it seems appropriate to close by briefly considering the merits of such arguments and the policy implications of further technological change.

Several headline-grabbing articles and books warn that this time may indeed be different and predict widespread displacement due to automation in the near future. For example, Frey & Osborne (2017) argued in a much cited article that as algorithms and robots can perform an increasingly wide range of non-routine cognitive and manual tasks, about 47 percent of US employment are at risk of automation in the next decade or two. Using a very similar approach, Frey et al. (2016) estimated the share of jobs at risk across the world, finding an average share of 57 percent in OECD countries and shares as high as 77 percent in China and 85 percent in Ethiopia. Authors like Brynjolfsson & McAfee (2014), MacCrory, Westerman, Alhammadi & Brynjolfsson (2016), and Elliott (2017) also foresee dramatic effects, with some skill groups facing displacement on a large scale. If such changes would come to pass, they would dwarf the essentially localised employment effects of RBTC.

On the other side, a number of researchers warn against what they perceive as automation hysteria. Arntz, Gregory & Zierahn (2016, 2017) argue that the occupation-centred approach of Frey & Osborne (2017) is susceptible to overestimates. Similarly, Autor (2015, p. 11), pointing to Polyani's paradox that "we know more than we can tell", and Acemoglu et al. (2014) maintain that humans are unlikely to soon be replaced by machines on a dramatic scale. Moreover, as Mokyr et al. (2015) point out, even if it would be possible to make informed predictions which occupations may suffer the greatest dislocation, it would remain impossible to predict which currently nonexistent occupations may emerge in the future. Thus, even large-scale displacement in some occupations may not necessarily lead to aggregate employment losses if displaced workers are absorbed by new sectors, as has occurred repeatedly throughout history.

My analyses do not suggest that widespread technological unemployment is around the corner. For example, there is no evidence for large-scale displacement of workers due to the proliferation of robots. Of course, the adoption of robots in production and especially services is still in its infancy and employment effects may yet become visible, and artificial intelligence may have a wider range of applicability. However, susceptibility does not mean that occupations will actually be automated. As we have seen, technology adoption depends on a myriad of factors, not least labour market institutions which may moderate the incentives for replacing workers (see also Wajcman 2017). The diversity of countries' experiences that has been central to the argument of this thesis clearly highlights the importance of the context in which the effects of technology play out. Therefore, the "alarmist" narratives in all likelihood overstate the negative impact of technology on aggregate employment, and I am inclined to side with the likes of Autor (2015), Arntz et al. (2016, 2017), and Mishel & Bivens (2017) who see no immediate danger of large-scale job losses due to technological change.

If this is true, it calls for a renewed focus on the quality of existing jobs. Far from holding a purely instrumental view of employment, people have over time come to place greater importance on the intrinsic qualities of their job (see, e.g., Gallie, Felstead & Green 2012 for Britain). This matters for the policy response to employment polarisation. One of the characteristics of employment polarisation is that workers who are displaced from their medium-skill, medium-wage job either have to retrain to find work in a higher-skilled occupation, or take a job below their skill level.[1] Large-scale polarisation would thus add to the existing problems of underemployment in terms of hours worked and utilisation of skills. Underemployment, however, is known to reduce wages and well-being (Bell & Blanchflower 2018, Heyes & Tomlinson 2021), so that workers affected by downward mobility in the wake of polarisation may face a double penalty in terms of the quality of their work, even as most of them are able to find new employment.

Perhaps, therefore, instead of an "end of work" debate, policymakers should focus on improving the quality of work, especially for those who have experienced job instability in the wake of technological change. Technology undoubtedly poses formidable challenges for employment and inequality, but there is as yet no evidence

[1] An alternative possibility is that such workers leave the labour force altogether. There is some evidence that especially older workers for whom retraining is less appealing have responded in this way to displacement due to technological change (Andrieu, Jamet, Marcolin & Squicciarini 2019).

that political processes cannot channel technological innovation in socially desirable ways. The task for social scientists is therefore not to decry the inevitability of the coming machine age, but to propose avenues how the outcomes that societies consider desirable could be attained.

Appendices

A
Appendix to chapter 4

A.1 2-digit ISCO-88 codes

11	Legislators and senior officials
12	Corporate managers
13	Managers of small enterprises
21	Physical, mathematical and engineering science professionals
22	Life science and health professionals
23	Teaching professionals
24	Other professionals
31	Physical and engineering science associate professionals
32	Life science and health associate professionals
33	Teaching associate professionals
34	Other associate professionals
41	Office clerks
42	Customer services clerks
51	Personal and protective services workers
52	Models, salespersons and demonstrators
61	Market-oriented skilled agricultural and fishery workers
62	Subsistence agricultural and fishery workers
71	Extraction and building trades workers

72	Metal, machinery and related trades workers
73	Precision, handicraft, printing and related trades workers
74	Other craft and related trades workers
81	Stationary plant and related operators
82	Machine operators and assemblers
83	Drivers and mobile plant operators
91	Sales and services elementary occupations
92	Agricultural, fishery and related labourers
93	Labourers in mining, construction, manfacturing and transport

B
Appendix to chapter 5

B.1 Methodological notes

This section discusses some additional methodological issues that have not been covered in detail in chapter 5.

B.1.1 Weights

I experimented with using the post-stratification weights provided in the EWCS dataset for calculating the indices but can confirm that the weights do not affect the occupational task intensity scores. For the calculation of my indices, weights would only be important if high- or low-routine or -complexity individuals would be systematically over- or under-sampled within an occupation. A comparison of a version of the index with and without weights (r = 0.9993) confirms that this is not the case. Hence, I do not use post-stratification weights to calculate the routine and complexity scores. Since the sample size is not identical in each country and wave, another potential concern is that countries or waves with larger samples could distort the findings. I checked for this possibility by calculating weights to give each country-wave the same weight in the combined indices. The correlation with the unweighted RTI index is r = 0.9988 and r = 0.9996 with the unweighted complexity index.

However, weights are relevant, and I indeed use them, for the calculation of occupational employment shares, as some occupations are over- or undersampled in national samples. The correlation between weighted and unweighted employment shares is only 0.8483. Thus, I include weights in the comparisons with other indices to account for the varying employment shares of occupations.

B.1.2 Distribution of individual responses

A final note concerns the distribution of individuals' responses to the survey questions included in the index. With most of the constituent questions being binary variables, and the answers being normalised to have a mean of 0 and a standard deviation of 1 in the relevant sample, the value for each individual depends on the overall distribution: the fewer individuals give the same answer, the further away from 0 will the value be. However, there is in all cases a substantial number of people answering either way, so that most values range between -1.5 and -0.5 and between 0.5 and 1.5. Since individuals' routine and complexity scores are the simple average of the questions making up the index, and it is rare for survey respondents to answer all questions affirmatively or negatively (the constituent questions are positively correlated, but not overly highly so), the individual-level scores in each occupation still have high standard deviations around 0.6.

The wide dispersion of individual answers does not change if I define the analytical groups more restrictively, such as occupation-sector combinations where one would expect tasks to be more similar than in broad occupations spanning all sectors. Partly, this is certainly due to measurement error which may be introduced by survey respondents misunderstanding the questions or understanding them differently, and various other complications. However, it above all illustrates that the tasks workers perform on their jobs vary considerably, and more than a single expert-assigned value would make one believe. Moreover, the survey is designed to minimise the issue of measurement error arising from overly subjective questions. In fact, the only two questions in my indices that require a truly subjective assessment are the questions about monotonous and complex tasks. Yet, when it comes to analysing the data, I nevertheless have to reduce this complexity to a single number.

B.2 Supplementary results

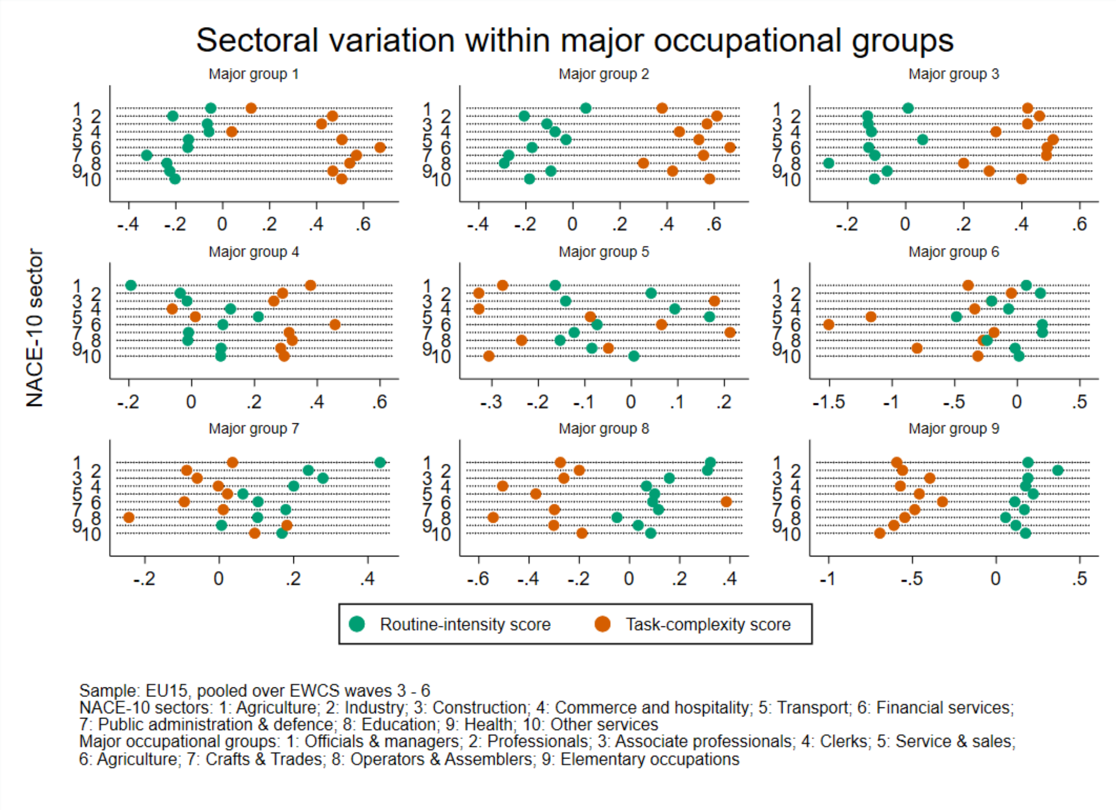

Figure B.1: *Sectoral variation within major occupational groups, EU-15 sample.*

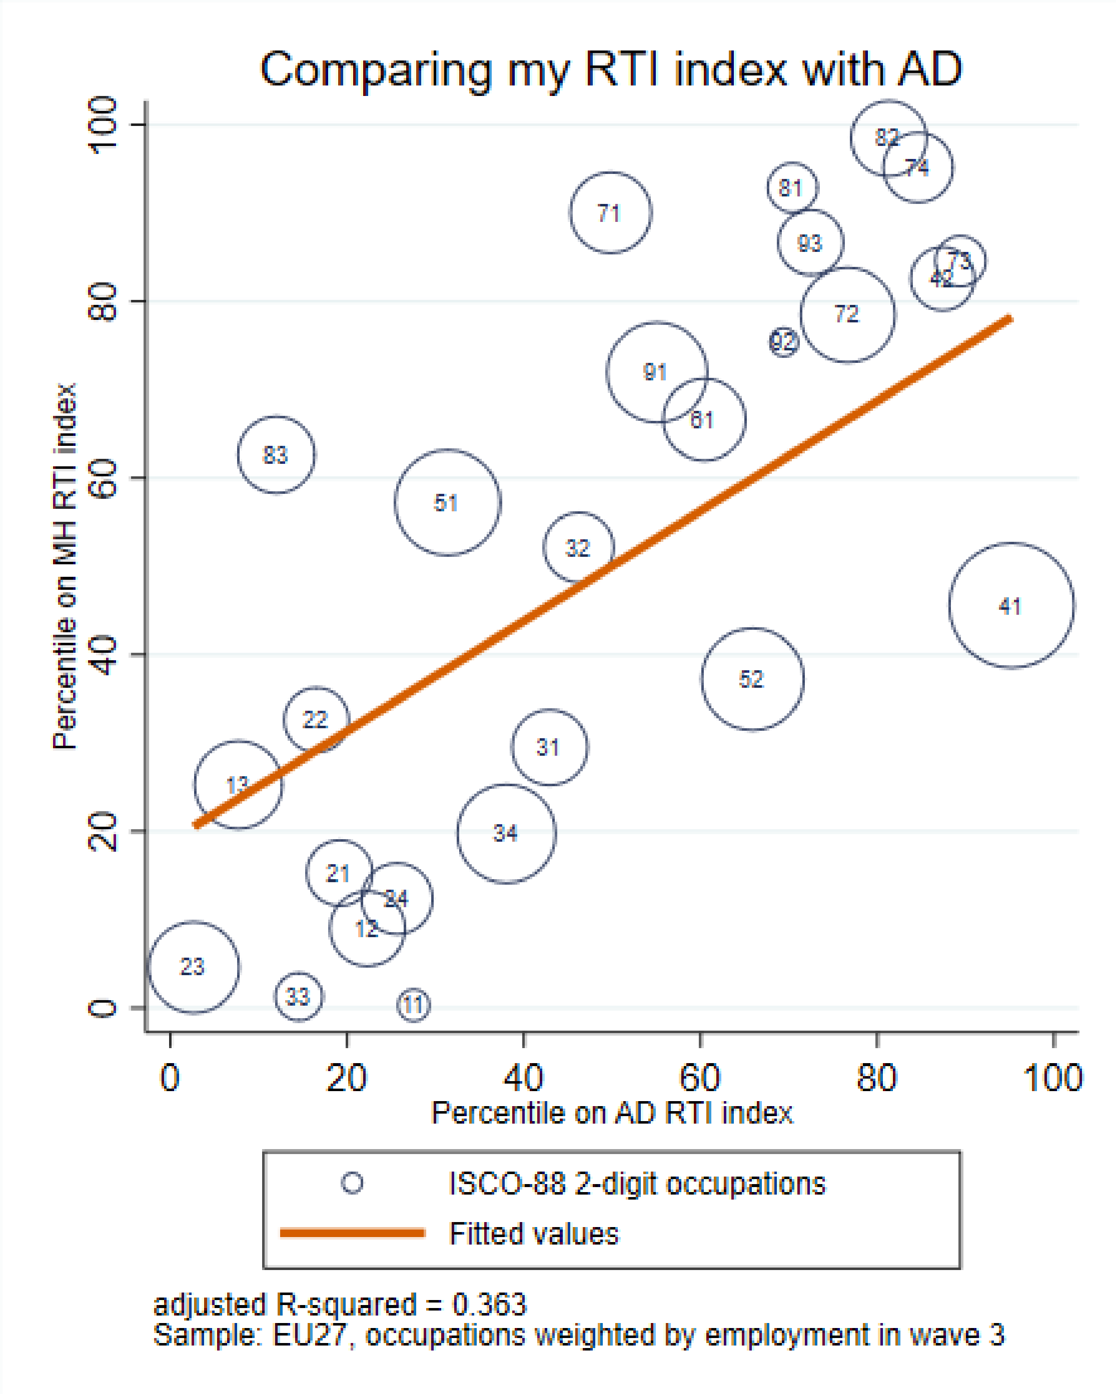

Figure B.2: *Comparison of my EU-27 routine index with Autor and Dorn 2013.*

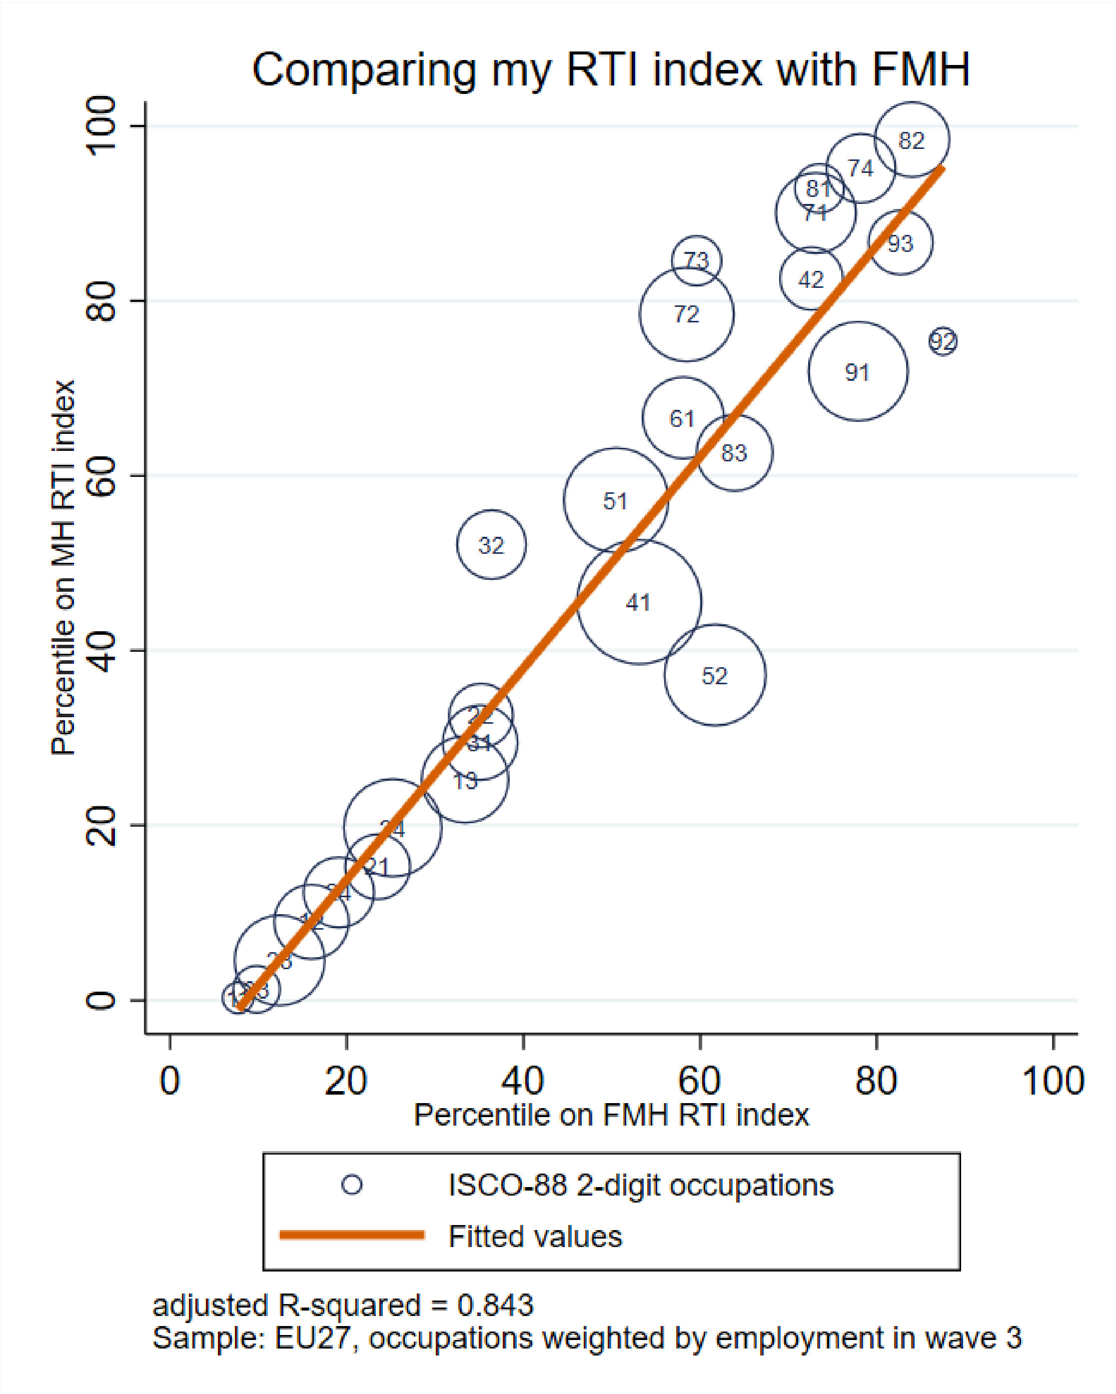

Figure B.3: *Comparison of my EU-27 routine index with Fernández-Macías and Hurley 2017.*

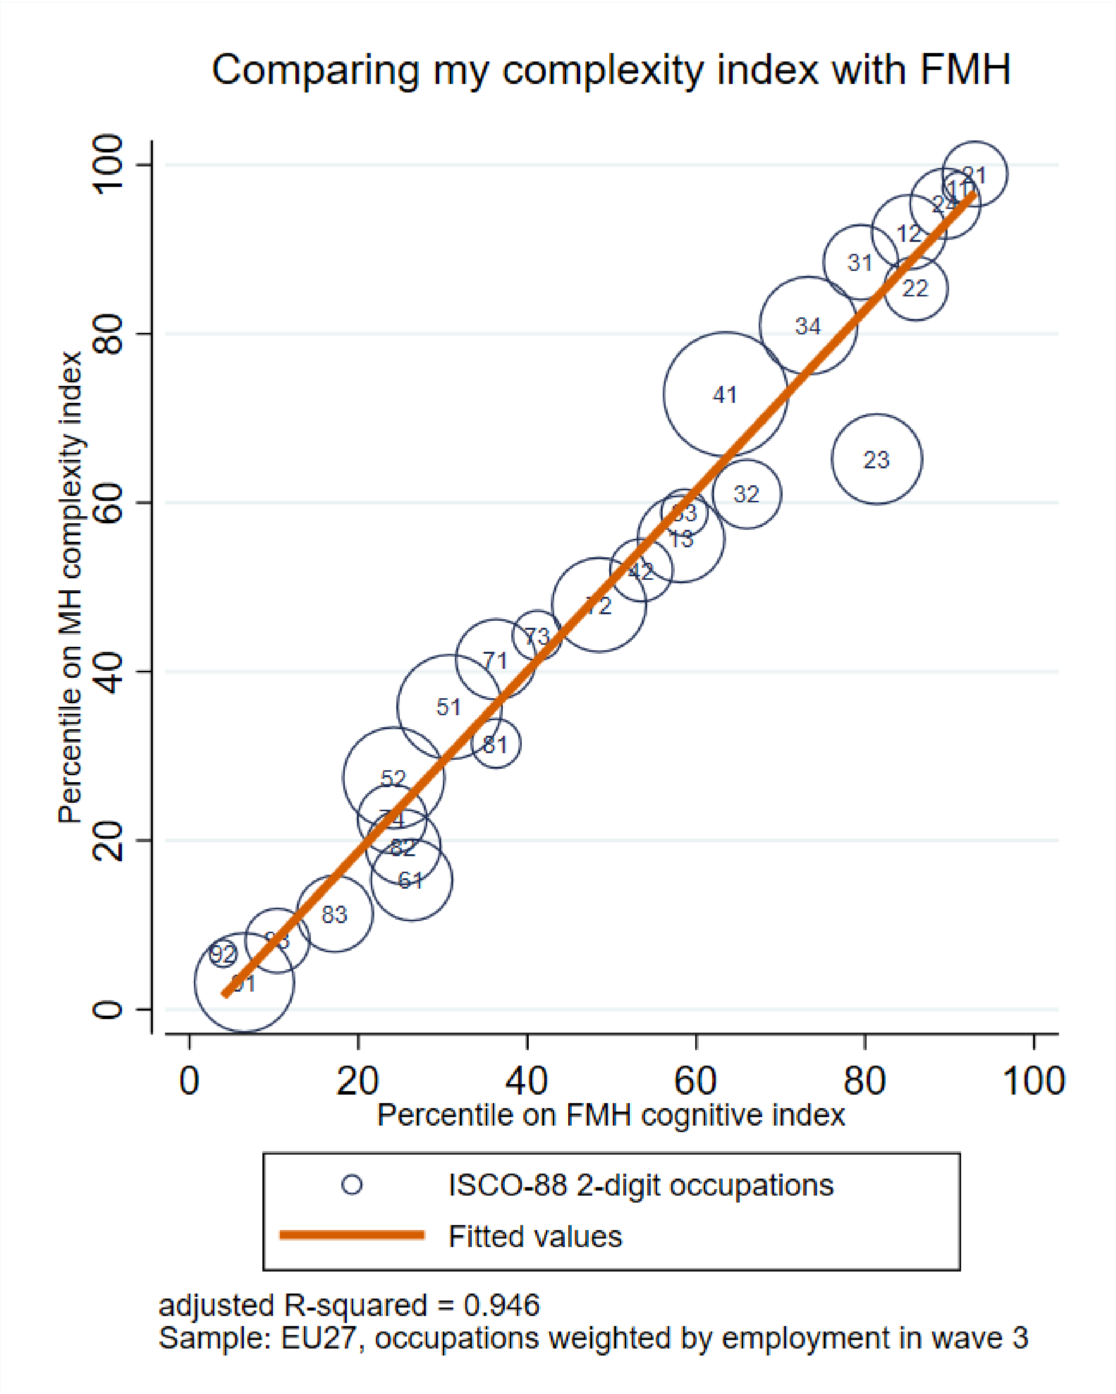

Figure B.4: *Comparison of my EU-27 complexity index with Fernández-Macías and Hurley 2017.*

C
Appendix to chapter 6

C.1 Supplementary figures

This appendix shows additional graphical evidence that has not been discussed in detail in various sections of chapter 6. Figure C.1 illustrates the - on average - upgrading nature of employment change by quintiles rather than terciles, as in chapter 6.3. Figure C.2 shows the smoothed employment changes in figure 6.3 by country, illustrating the diversity of country patterns. Figure C.3 displays the employment changes by wage tercile restricted to the 1995 - 2013 time frame, while figures C.4, C.5, C.6, and C.7 show RTI-wage curves using different wage measures and worker populations.

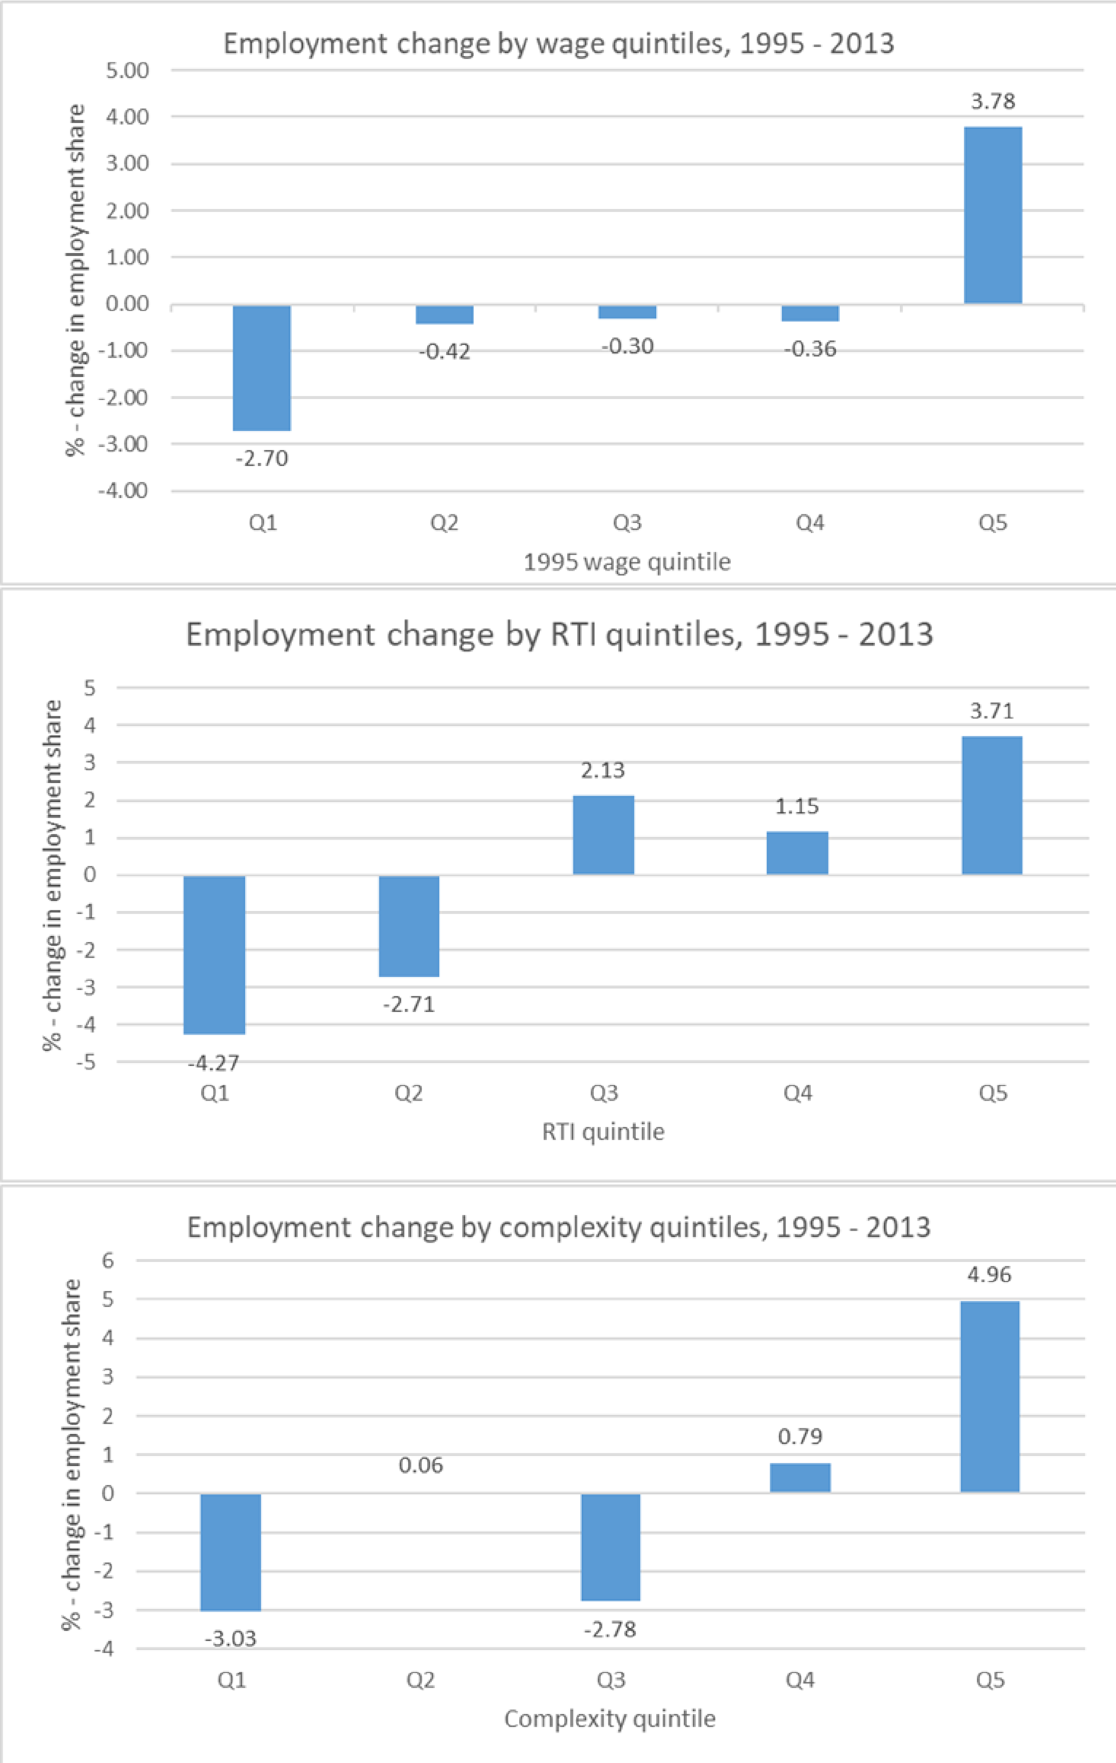

Figure C.1: *Employment changes by wage, RTI, and complexity quintiles, 1995 - 2013.*

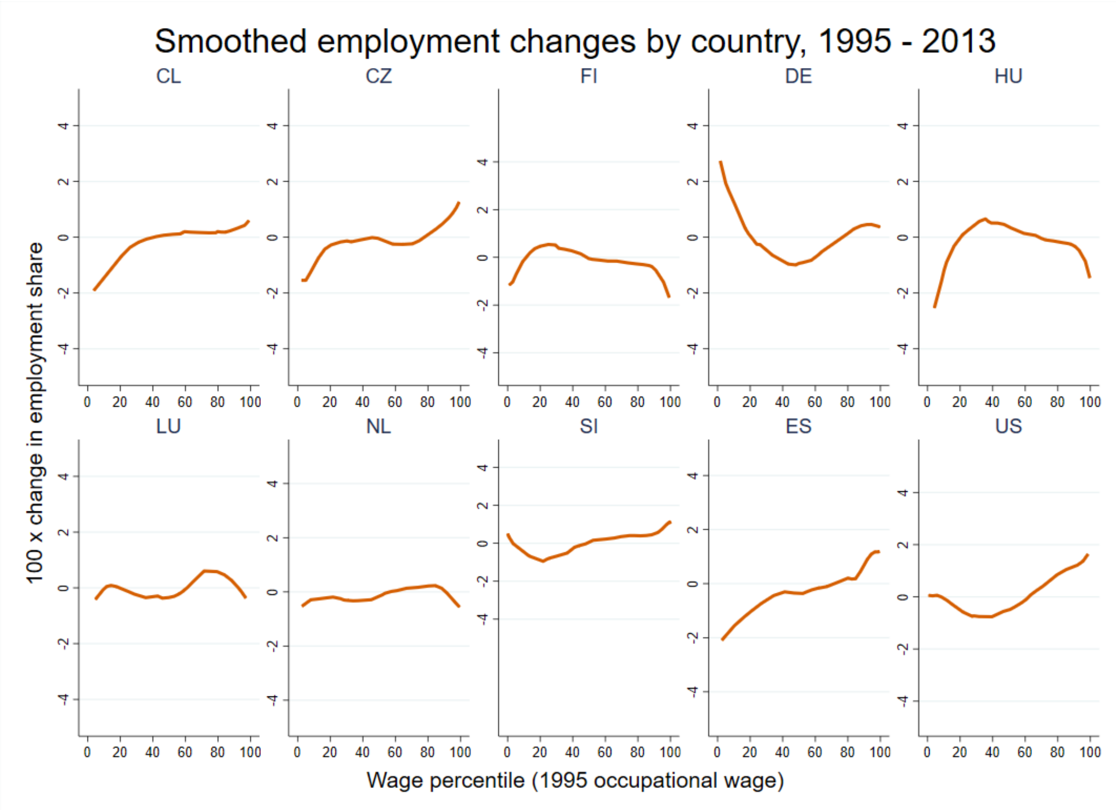

Figure C.2: *Smoothed employment changes in individual countries, 1995 - 2013. Employment changes are estimated using a locally smoothed regression with bandwidth = 0.75.*

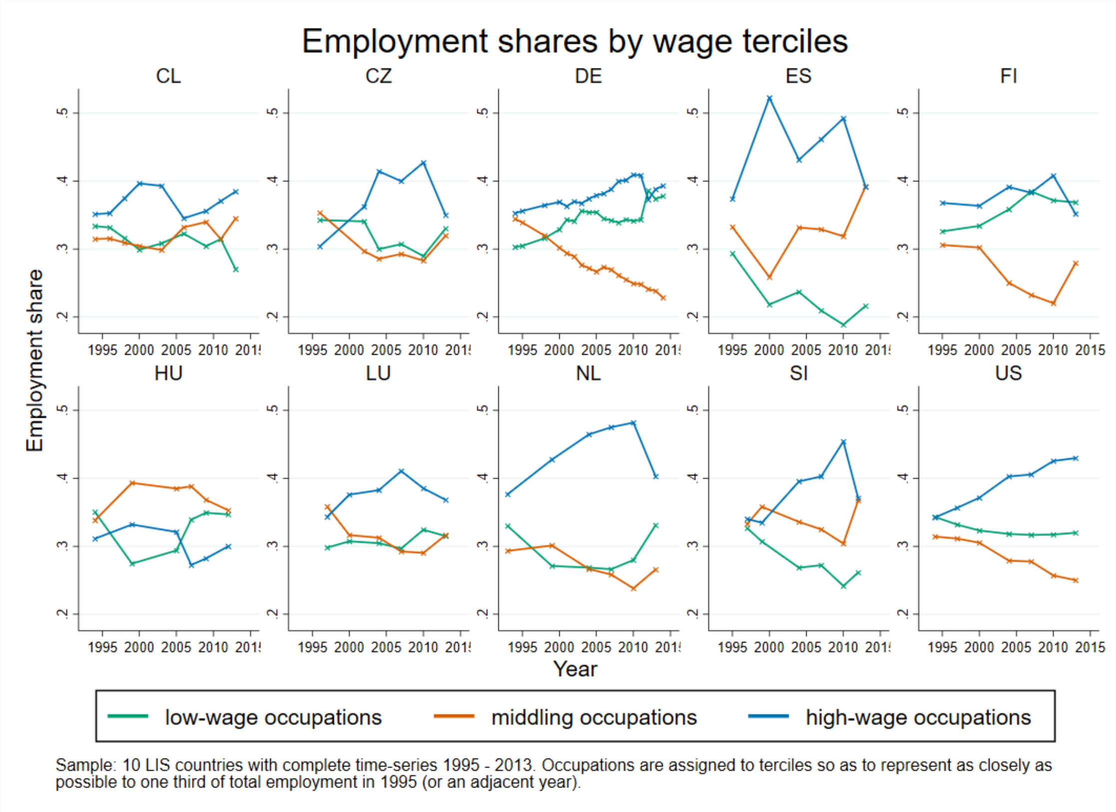

Figure C.3: *Employment change by wage tercile between 1995 and 2013. Compared to figure 6.2, in some cases occupations are assigned to different wage terciles based on the different initial period.*

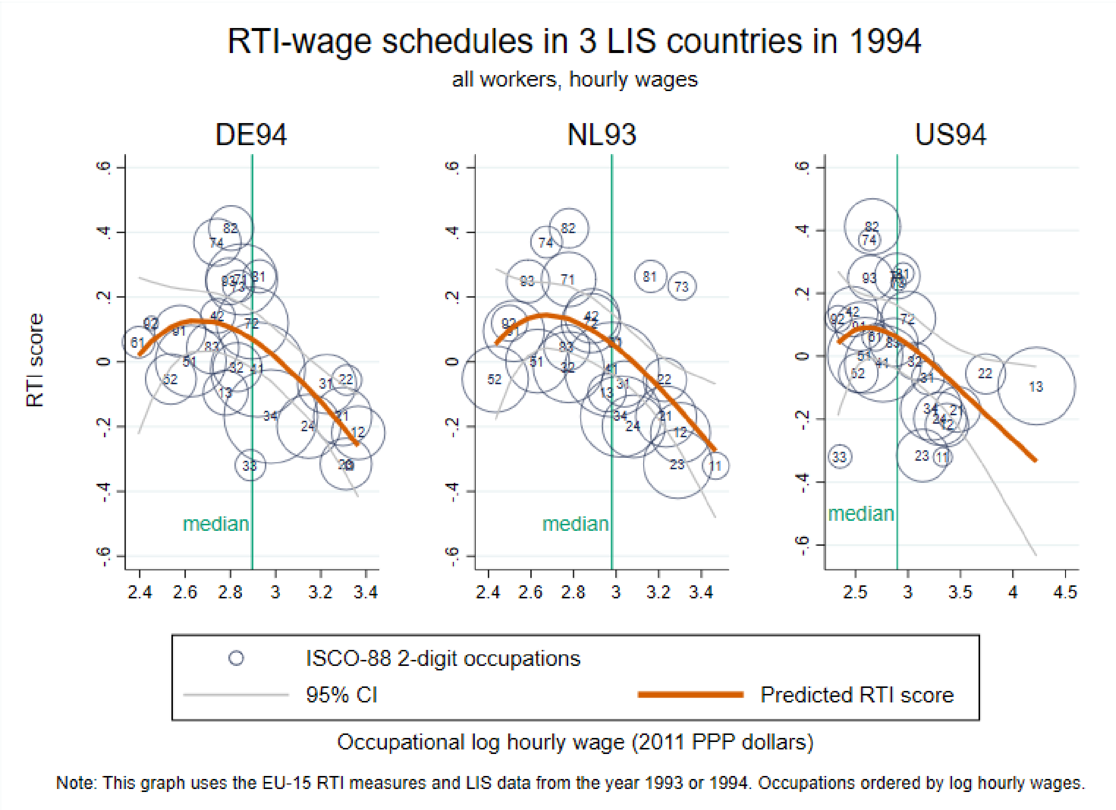

Figure C.4: *RTI-wage schedules with hourly wage data, at the beginning of the period of analysis (1994).*

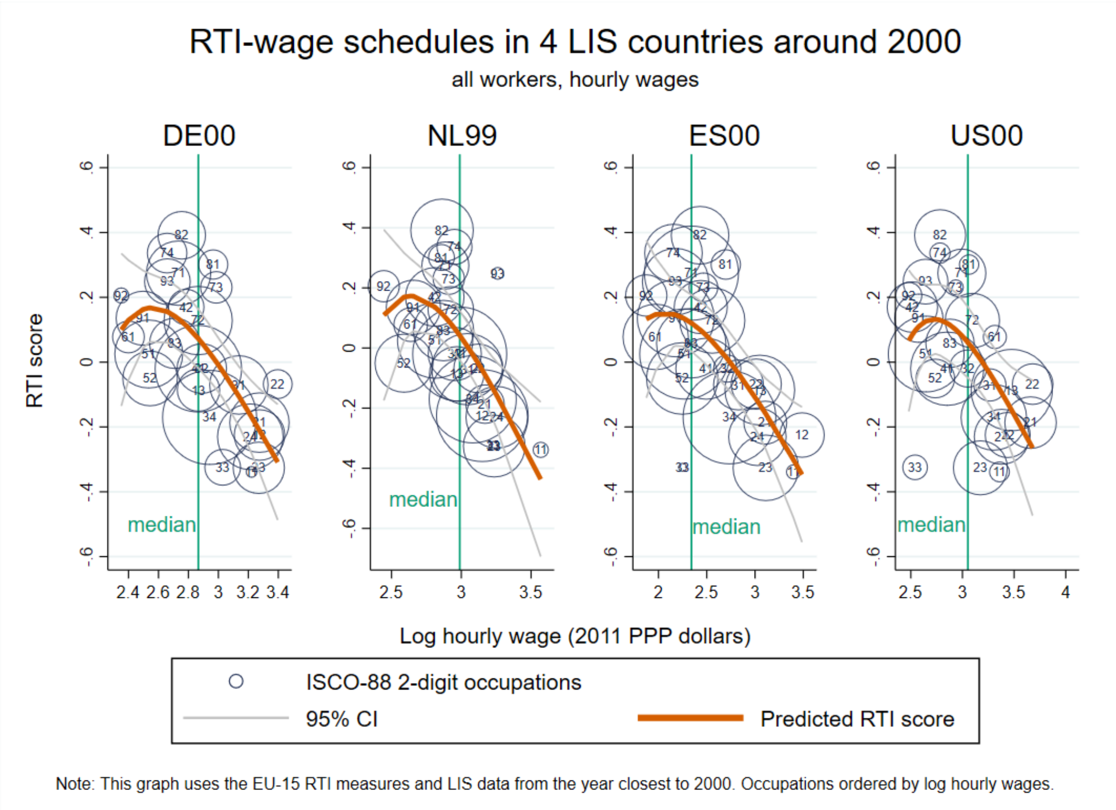

Figure C.5: *RTI-wage schedules with hourly wage data, 2000. The figures here are identical to the hourly wage figures in figure 6.6, only the aspect ratio is different.*

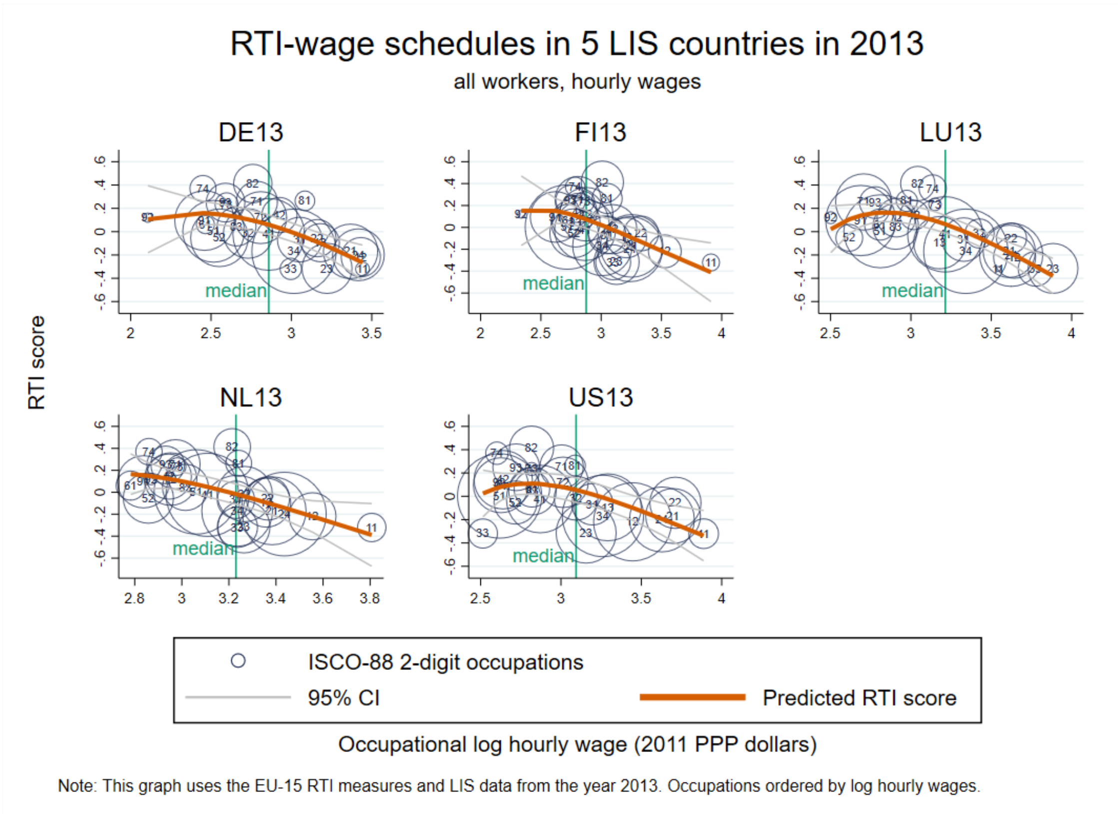

Figure C.6: *RTI-wage schedules with hourly wage data, at the end of the period of analysis (2013).*

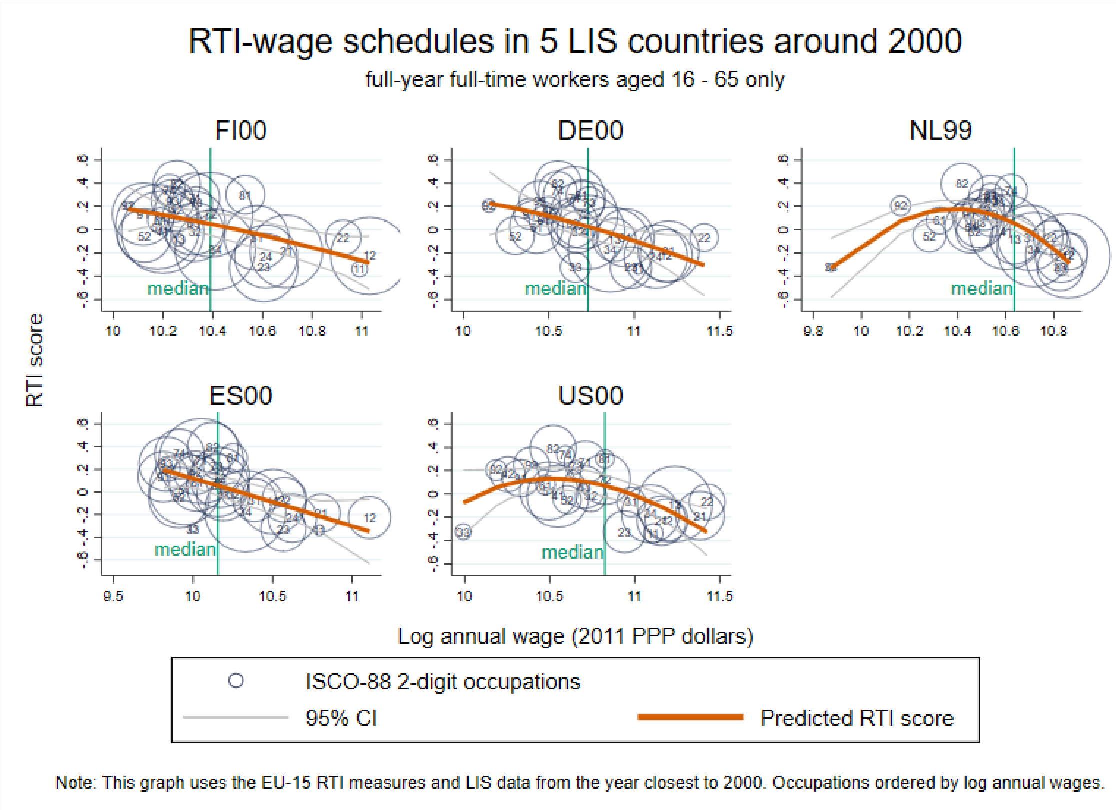

Figure C.7: *RTI-wage schedules using annual wages of FYFT workers in countries with information on worker status in 2000.*

C.2 Supplementary analyses

This appendix shows additional regression evidence and descriptive analyses that have not been discussed in detail in chapter 6.

Table C.1: Occupational wages and employment change, 1995 - 2013

PANEL A: FULL SAMPLE				
	All workers		Women only	Men only
DV: ΔEmployment '95-'13	(A1)	(A2)	(A3)	(A4)
1995 log annual wage	0.107**	0.008*	0.016	0.139**
	(0.047)	(0.003)	(0.009)	(0.05)
1995 squared log wage	-0.005*			-0.007**
	(0.002)			(0.003)
Observations	260	260	252	260
R-squared	0.061	0.052	0.067	0.071
PANEL B: CONTROLLING FOR FYFT SHARE				
	All workers		Women only	Men only
DV: ΔEmployment '95-'13	(B1)	(B2)	(B3)	(B4)
1995 log annual wage	0.083	0.024**	0.026	0.172**
	(0.066)	(0.006)	(0.014)	(0.057)
1995 squared log wage	-0.003			-0.007*
	(0.003)			(0.003)
1995 FYFT share	-0.070**	-0.069**	-0.073	-0.055*
	(0.02)	(0.019)	(0.038)	(0.022)
Observations	130	130	127	130
R-squared	0.093	0.091	0.072	0.094

Standard errors clustered by country in parentheses. All occupations are weighted by their initial employment share in 1995. All regressions include a constant and country-dummies. Sample: 10-country sample (panel A)/DE, ES, FI, NL, US (panel B).*** p<0.01, ** p<0.05, * p<0.1

Table C.2: Tasks and employment demand, additional analyses

PANEL A: BY WORKER STATUS						
	ALL WORKERS			FYFT WORKERS ONLY		
DV: Employment share	(A1)	(A2)	(A3)	(A4)	(A5)	(A6)
Trend interacted with:						
RTI	-0.604***		-0.838***	-0.432*		-0.590**
	(0.200)		(0.267)	(0.231)		(0.296)
Complexity		0.133	-0.182*		0.082	-0.125
		(0.083)	(0.103)		(0.123)	(0.155)
Observations	2,930	2,930	2,930	2,930	2,930	2,930
Country-occupations	597	597	597	597	597	597
R-squared	0.053	0.025	0.059	0.025	0.019	0.026
PANEL B: BY SEX						
	ALL WOMEN			ALL MEN		
DV: Employment share	(B1)	(B2)	(B3)	(B4)	(B5)	(B6)
Trend interacted with:						
RTI	-0.858***		-1.190***	-0.457**		-0.559**
	(0.311)		(0.447)	(0.203)		(0.273)
Complexity		0.158	-0.252		0.142	-0.078
		(0.141)	(0.204)		(0.091)	(0.117)
Observations	2,891	2,891	2,891	2,923	2,923	2,923
Country-occupations	597	597	597	597	597	597
R-squared	0.063	0.038	0.068	0.042	0.032	0.043
PANEL C: ACCOUNTING FOR NON-STANDARD EMPLOYMENT						
DV: Employment share	(C1)	(C2)	(C3)			
FYFT share	-0.007	-0.008	-0.007			
	(0.006)	(0.006)	(0.006)			
Trend interacted with:						
RTI	-0.774***		-0.974***			
	(0.219)		(0.285)			
Complexity		0.183	-0.149			
		(0.121)	(0.149)			
Observations	933	933	933			
Country-occupations	182	182	182			
R-squared	0.068	0.031	0.072			

All point estimates and standard errors are multiplied by 100. Robust standard errors are clustered by country-occupation. All occupations are weighted by their initial employment share in 1995. All regressions include country-occupation fixed-effects and wave fixed-effects. Sample: all LIS countries 1995 – 2016 (panels A and B), all countries with FYFT data 1995 – 2016 (panel C). *** p<0.01, ** p<0.05, * p<0.1

Table C.3: Routine intensity by wage tercile and country

Country-year	Low-wage	Medium-wage	High-wage	Country average
CL ’94	0.21	0.11	-0.08	0.07
CZ ’96	0.15	0.03	-0.01	0.05
DE ’94	0.12	0.09	-0.18	-0.00
ES ’95	0.22	0.07	-0.08	0.09
FI ’95	0.10	0.10	-0.17	-0.01
HU ’94	0.14	0.13	-0.18	0.04
LU ’97	0.09	0.08	-0.20	-0.01
NL ’93	0.06	0.02	-0.20	-0.02
SI ’97	0.25	0.06	-0.15	0.05
US ’94	0.08	0.05	-0.09	0.00

This table shows the employment-weighted average routine-intensity scores for low-, medium-, and high-wage occupations in each country at the beginning of the period of analysis. Occupational routine scores do not differ between countries, hence differences are due to different employment shares and allocations of occupations to wage terciles.

C.3 RTI-wage curves over time in individual countries

Figure C.8 shows the RTI-wage curves in all available country-years for a sample of 16 countries.

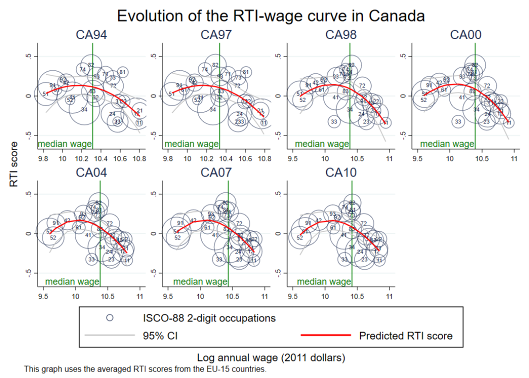
Evolution of the RTI-wage curve in Canada
CA94 CA97 CA98 CA00
CA04 CA07 CA10
median wage
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

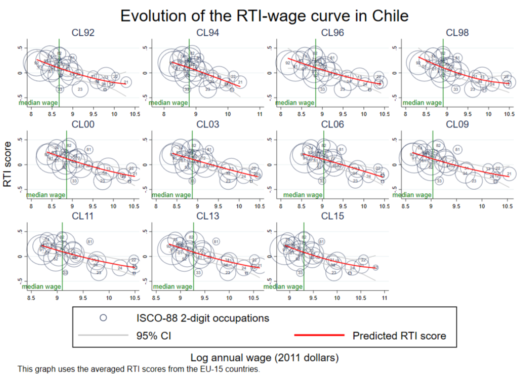
Evolution of the RTI-wage curve in Chile
CL92 CL94 CL96 CL98
CL00 CL03 CL06 CL09
CL11 CL13 CL15
median wage
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

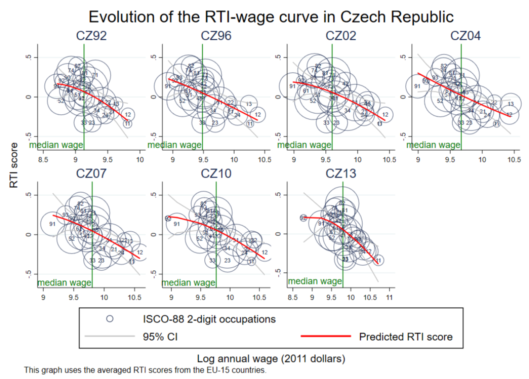
Evolution of the RTI-wage curve in Czech Republic
CZ92 CZ96 CZ02 CZ04
CZ07 CZ10 CZ13
median wage
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

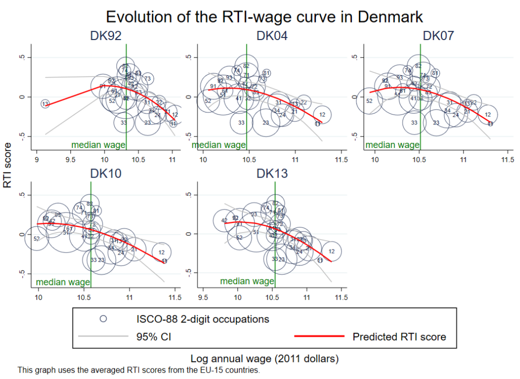
Evolution of the RTI-wage curve in Denmark
DK92 DK04 DK07
DK10 DK13
median wage
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

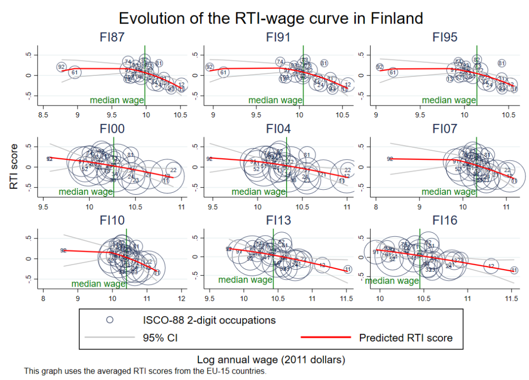
Evolution of the RTI-wage curve in Finland
FI87 FI91 FI95
FI00 FI04 FI07
FI10 FI13 FI16
median wage
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

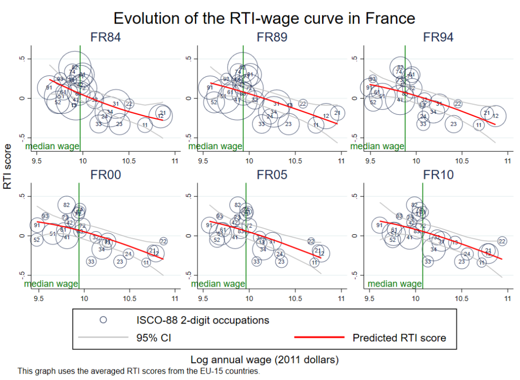
Evolution of the RTI-wage curve in France
FR84 FR89 FR94
FR00 FR05 FR10
median wage
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

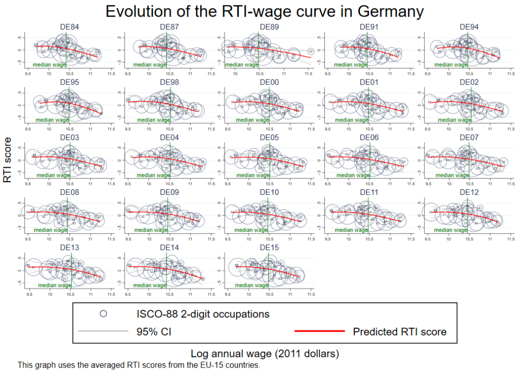
Evolution of the RTI-wage curve in Germany
DE84 DE87 DE89 DE91 DE94
DE95 DE98 DE00 DE01 DE02
DE03 DE04 DE05 DE06 DE07
DE08 DE09 DE10 DE11 DE12
DE13 DE14 DE15
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

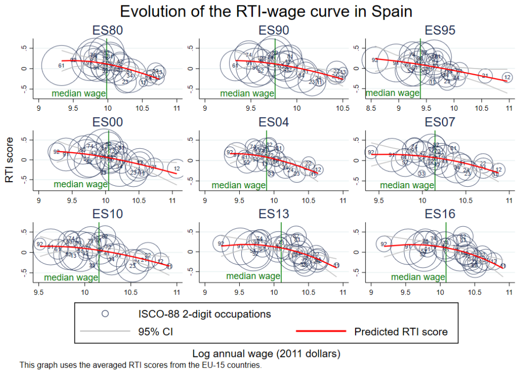
Evolution of the RTI-wage curve in Spain
ES80 ES90 ES95
ES00 ES04 ES07
ES10 ES13 ES16
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

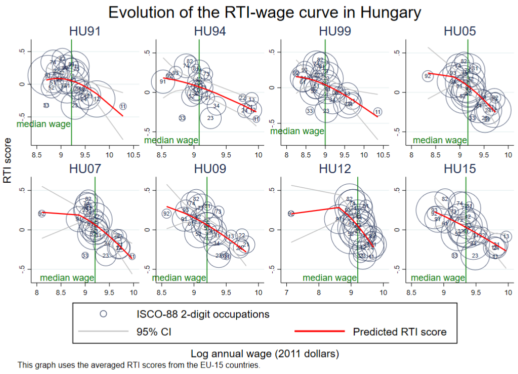
Evolution of the RTI-wage curve in Hungary
HU91 HU94 HU99 HU05
HU07 HU09 HU12 HU15
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

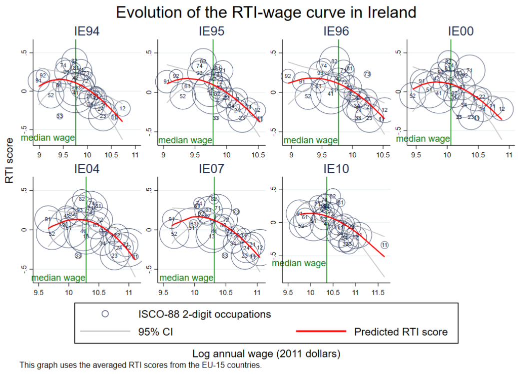
Evolution of the RTI-wage curve in Ireland
IE94 IE95 IE96 IE00
IE04 IE07 IE10
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

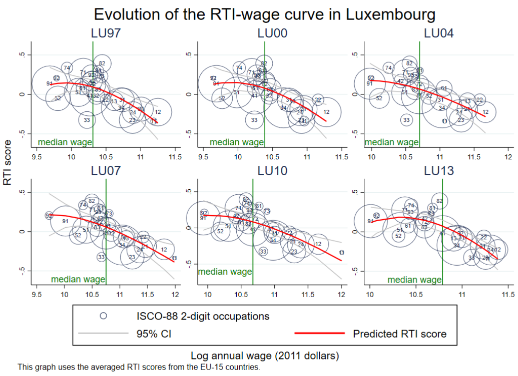
Evolution of the RTI-wage curve in Luxembourg
LU97 LU00 LU04
LU07 LU10 LU13
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

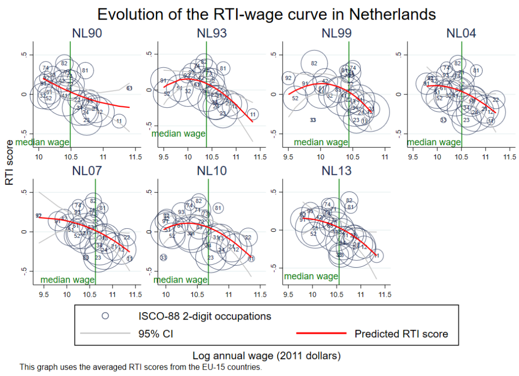
Evolution of the RTI-wage curve in Netherlands
NL90 NL93 NL99 NL04
NL07 NL10 NL13
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

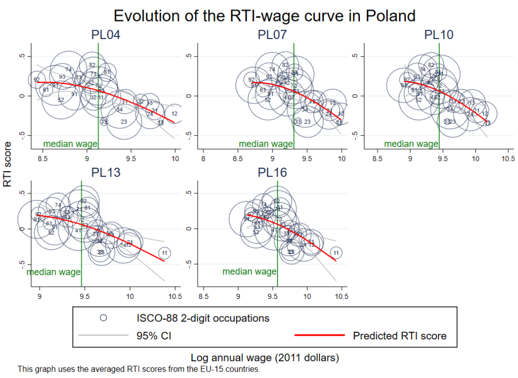
Evolution of the RTI-wage curve in Poland
PL04 PL07 PL10
PL13 PL16
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

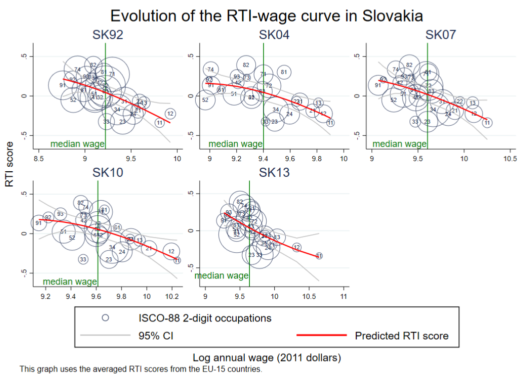
Evolution of the RTI-wage curve in Slovakia
SK92 SK04 SK07
SK10 SK13
RTI score
ISCO-88 2-digit occupations
95% CI
Predicted RTI score
Log annual wage (2011 dollars)
This graph uses the averaged RTI scores from the EU-15 countries.

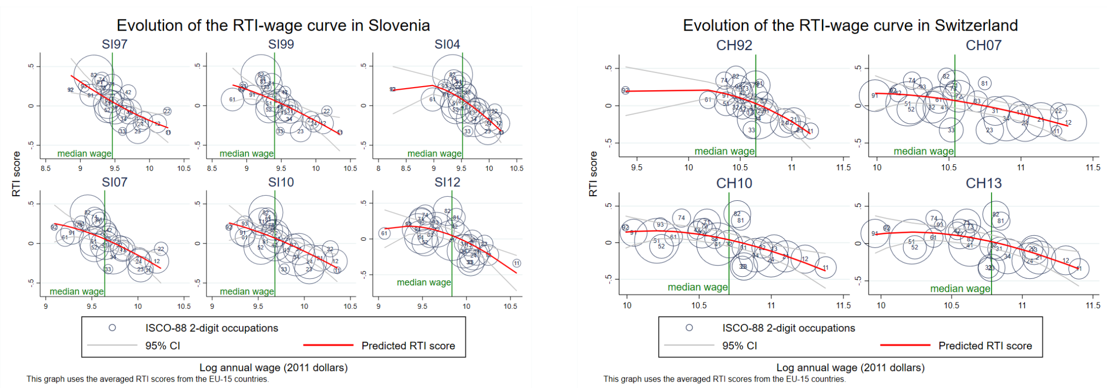

Figure C.8: RTI-wage curves over time in individual countries.

C.4 Changes to the wage structure

This section expands upon the analysis of changes in occupational wage rankings in chapter 6.4.4. In figure C.9, I provide a full overview of how occupations are distributed across and have moved between wage percentiles in individual countries between 1995 and 2013. The distribution matrix illustrates that the ordering of the ISCO classification only partly corresponds to the occupational wage hierarchy. The occupations in the highest and lowest ISCO groups are also the best and worst paid, respectively. Conversely, in ISCO groups 4 through 8, the ordering of wages does not follow the hierarchy according to ISCO. Some occupations in major groups 7 and especially 8 are quite well paid, whereas major groups 5 and 6 have very low wages.

However, the key outcome of interest are changes in the occupational wage hierarchy. What is apparent, is that professionals and health and teaching associate professionals (groups 32 and 33) have generally improved their positions as well as, interestingly, plant and machine operators (groups 81 and 82). This shows that even shrinking occupations can achieve significant wage growth. The biggest losers have been most trades workers (groups 71, 72, 73), drivers (group 83), as well as some professional and associate professional groups (teachers and physical/engineering associate professionals). Thus, among the winners and losers are both high- and low-routine occupations, and even occupations of the same major group (groups 32 and 33, versus groups 31 and 34). Clearly, routine-intensity alone cannot explain the reshuffling of the wage hierarchy that has taken place between 1995 and 2013. There is a fruitful avenue for future research in zooming in on these occupations and investigate which factors, beyond their routine-intensity, have affected their wage trajectory. Here, one has to account for the fact that the employment expansion of the top occupations necessarily exerts a downward pressure on the wage percentiles of lower-ranking occupations.

Looking for differences between countries, no systematic differences are apparent between countries with a hump-shaped RTI-wage curve and countries without. The patterns appear to be broadly similar across all countries. The intensity of changes varies quite widely between countries but can likely be attributed to the different sample sizes (Chile, Germany, and the US have the largest samples and the most stable wage structures as measured by the absolute sum of changes). The differences in data quality prevent all too far-reaching interpretations of the data presented here. Nevertheless, the figure illustrates neatly that there are broadly similar patterns across countries which occupations have moved up or down the wage hierarchy.

	1995										2013										Change										
ISCO code	CL	CZ	FI	DE	HU	LU	NL	SI	ES	US	CL	CZ	FI	DE	HU	LU	NL	SI	ES	US	CL	CZ	FI	DE	HU	LU	NL	SI	ES	US	∑ of changes
11	97	98	99	100	100	94	100	100	62	80	95	100	100	100	94	89	99	100	100	99	-2	2	1	0	-6	-5	-1	0	38	19	46
12	97	99	96	98	97	97	97	98	100	92	95	98	98	98	81	98	97	98	99	80	-2	-1	2	0	-16	1	0	0	-1	-12	-29
13	92	97	5	56	93	64	80	85	90	85	89	88	86	90	94	55	76	74	57	72	-3	-9	81	34	1	-9	-4	-11	-33	-13	34
21	99	93	83	94	91	92	92	95	99	95	98	95	92	94	98	81	92	93	96	93	-1	2	9	0	7	-11	0	-2	-3	-2	-1
22	98	94	62	91	86	63	61	100	92	99	96	86	96	82	95	84	73	96	90	97	-2	-8	34	-9	9	21	12	-4	-2	-2	49
23	82	85	91	83	77	84	85	89	88	63	86	69	75	78	74	87	52	80	87	54	4	-16	-16	-5	-3	3	-33	-9	-1	-9	-85
24	96	95	74	88	88	88	74	93	94	77	93	92	81	86	92	93	83	88	93	88	-3	-3	7	-2	4	5	9	-5	-1	11	22
31	88	89	79	79	72	79	78	82	97	67	83	83	63	73	78	74	68	64	74	58	-5	-6	-16	-6	6	-5	-10	-18	-23	-9	-92
32	79	46	39	21	37	53	22	75	54	54	76	79	69	27	56	72	57	84	77	46	-3	33	30	6	19	19	35	9	23	-8	163
33	41	49	86	30	12	29		66		0	56	71	66	31	72	77	49	76	80	1	15	22	-20	1	60	48		10		1	137
34	86	81	67	71	83	72	67	70	85	71	80	75	55	65	86	64	63	69	83	64	-6	-6	-12	-6	3	-8	-4	-1	-2	-7	-49
41	76	39	45	37	43	36	45	60	58	40	60	50	37	39	50	48	43	53	63	33	-16	11	-8	2	7	12	-2	-7	5	-7	-3
42	59	59	29	8	47	46	31	78	41	3	46	25	29	48	65	38	22	59	47	14	-13	-34	0	40	18	-8	-9	-19	6	11	-8
51	36	26	25	12	32	18	14	37	34	22	40	15	16	11	39	16	15	15	18	6	4	-11	-9	-1	7	-2	1	-22	-16	-16	-65
52	45	13	17	5	26	12	8	30	21	12	32	9	7	17	17	10	8	24	26	26	-13	-4	-10	12	-9	-2	0	-6	5	14	-13
61	22	17	3	19	39	21	29	41	16	27	25	24	26	34	22	27	26	0	39	28	3	7	23	15	-17	6	-3	-41	23	1	17
71	53	75	20	50	51	25	37	48	67	47	52	29	32	56	31	23	32	9	36	43	-1	-46	12	6	-20	-2	-5	-39	-31	-4	-130
72	63	65	56	61	59	50	57	53	75	58	66	60	49	52	68	36	35	38	53	49	3	-5	-7	-9	9	-14	-22	-15	-22	-9	-91
73	39	71	52	46	70	59	52	26	62	49	23	22	41	21	45	39	24	34	50	39	-16	-49	-11	-25	-25	-20	-28	8	-12	-10	-188
74	31	21	14	44	16	14	28	24	27	28	18	21	27	30	25	30	25	28	33	21	-13	0	13	-14	9	16	-3	4	6	-7	11
81	71	33	87	54	36	58	89	22	82	61	77	66	72	80	83	40	71	62	72	57	6	33	-15	26	47	-18	-18	40	-10	-4	87
82	57	30	50	24	21	61	34	13	52	32	48	42	60	46	10	41	70	45	70	37	-9	12	10	22	-11	-20	36	32	18	5	95
83	68	54	35	27	66	56	54	44	80	51	72	35	44	32	61	33	28	31	42	40	4	-19	9	5	-5	-23	-26	-13	-38	-11	-117
91	14	2	9	2	4	5	3	3	10	9	10	2	2	4	3	5	2	4	7	19	-4	0	-7	2	-1	0	-1	1	-3	10	-3
92	4	5	1	7	9	9	17	0	3	6	2	0	0	0	0	9	4	1	1	21	-2	-5	-1	-7	-9	0	-13	1	-2	15	-23
93	27	8	31	17	10	43	26	1	47	15	21	5	24	23	6	29	5	20	31	23	-6	-3	-7	6	-4	-14	-21	19	-16	8	-38
																				Absolute sum of changes	159	347	370	261	332	292	296	336	340	225	

Figure C.9: *Changes in occupations' average wage rank.*

D
Appendix to chapter 7

Table D.1: Cross-classification of sector and occupation, based on EWCS data from the EU-15

ISCO-88 2-digit	Agriculture %	Industry %	Services %	Total N
11	0.39	4.31	95.29	255
12	2.92	27.90	69.17	2,394
13	3.00	15.30	81.70	3,928
21	0.76	30.46	68.79	2,108
22	0.28	2.16	97.57	2,178
23	0.07	0.32	99.61	4,114
24	0.20	11.16	88.65	3,021
31	0.45	38.55	61.00	2,236
32	0.47	3.73	95.79	2,330
33	0.08	0.79	99.13	1,260
34	0.39	12.63	86.98	5,944
41	0.78	17.87	81.35	7,661
42	0.08	5.80	94.12	2,535
51	0.27	1.74	98.00	8,631
52	0.48	7.06	92.46	4,986
61	86.41	1.70	11.89	2,002
62	96.44	0.44	3.11	225
71	0.32	88.85	10.83	3,481
72	0.44	58.52	41.03	3,385
73	0.36	72.69	26.95	835
74	0.41	59.29	40.30	1,938
81	0.72	85.43	13.85	693
82	1.18	86.46	12.36	1,942
83	2.25	18.17	79.57	2,619
91	0.89	6.16	92.95	6,596
92	80.28	5.05	14.68	436
93	0.92	55.05	44.02	1,840
Total	**3.61**	**21.49**	**74.9**	**79,573**

Table D.2: Overview of country-years included in the main regression sample

Country/Year	93	94	95	96	97	98	99	00	01	02	03	04	05	06	07	08	09	10	11	12	13	14	15	16
Austria												✓			✓			✓			✓			
Belgium			✓		✓			✓			✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	
Canada		✓			✓	✓		✓				✓			✓			✓						
Czechia				✓						✓		✓			✓			✓			✓			
Denmark												✓			✓			✓			✓			
Estonia																		✓			✓			
Finland			✓					✓				✓			✓			✓			✓			
France		✓						✓					✓					✓						
Germany		✓	✓			✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	
Greece												✓			✓			✓			✓			
Hungary		✓					✓						✓		✓		✓			✓			✓	
Iceland																		✓						
Ireland		✓	✓	✓				✓				✓			✓			✓	✓	✓	✓	✓	✓	
Netherlands	✓						✓					✓			✓			✓			✓			
Poland												✓			✓			✓			✓			
Slovakia												✓			✓			✓			✓	✓	✓	
Slovenia																		✓		✓				
Spain			✓					✓				✓			✓			✓			✓			
Switzerland														✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	
United States	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Note: N = 150 country-years. Belgium, Canada, Denmark, Germany, Iceland, Ireland, The Netherlands, and the United States are coded as exhibiting a hump-shaped routine-wage curve, the remaining countries predominantly have a monotonic RTI-wage curve.

Table D.3: Correlations in the main regression sample

Variables	**(1)**	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**	**(8)**
(1) Wage ratio	1							
(2) Robot density	0.112	1						
(3) Union density	-0.234*	0.035	1					
(4) EPL (temporary)	-0.445*	0.157	0.193	1				
(5) Trade openness	-0.438*	-0.143	0.272*	0.146	1			
(6) Capital openness	0.197	0.320*	-0.069	0.034	-0.08	1		
(7) Manufac. emp.	-0.300*	-0.104	-0.07	0.195	0.16	-0.371*	1	
(8) RTI average	-0.385*	-0.431*	-0.15	0.264*	0.101	-0.401*	0.683*	1
(9) Real GDP growth	0.108	-0.262*	0.038	-0.175	0.158	-0.065	0.121	0.143

* $p<0.01$

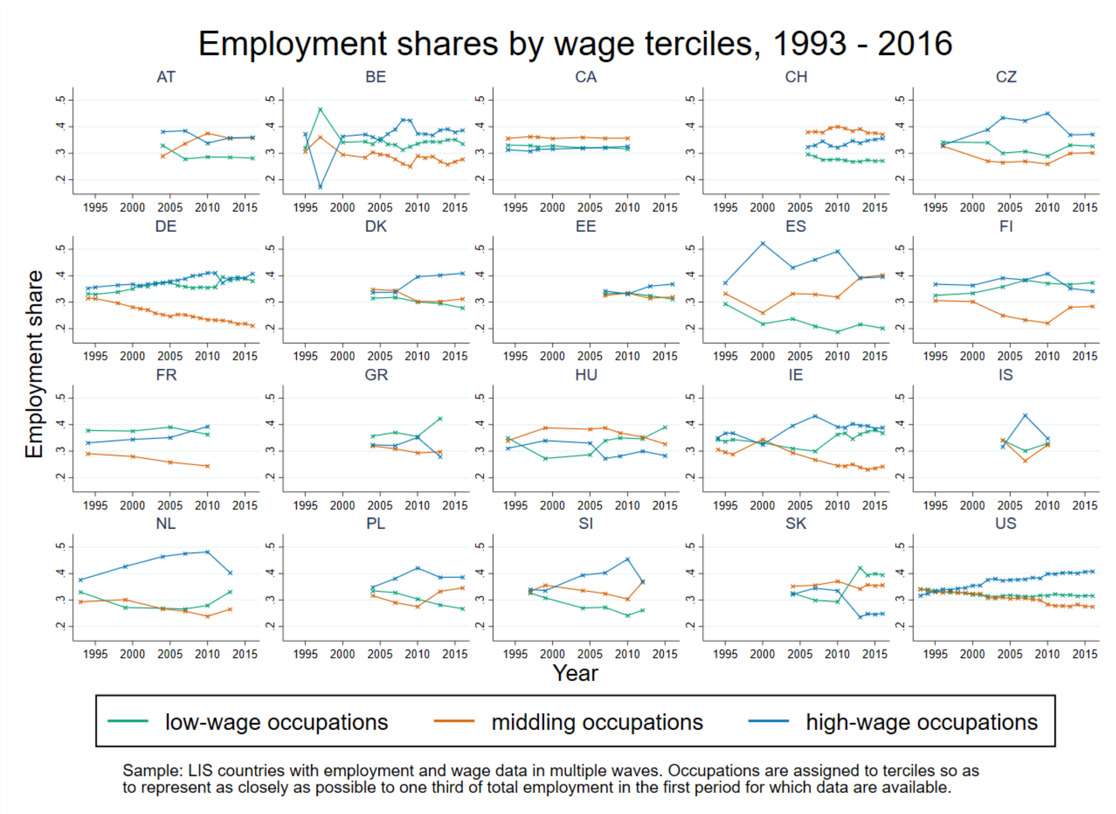

Figure D.1: *Employment shares by wage tercile in 20 OECD countries, 1993 - 2016.*

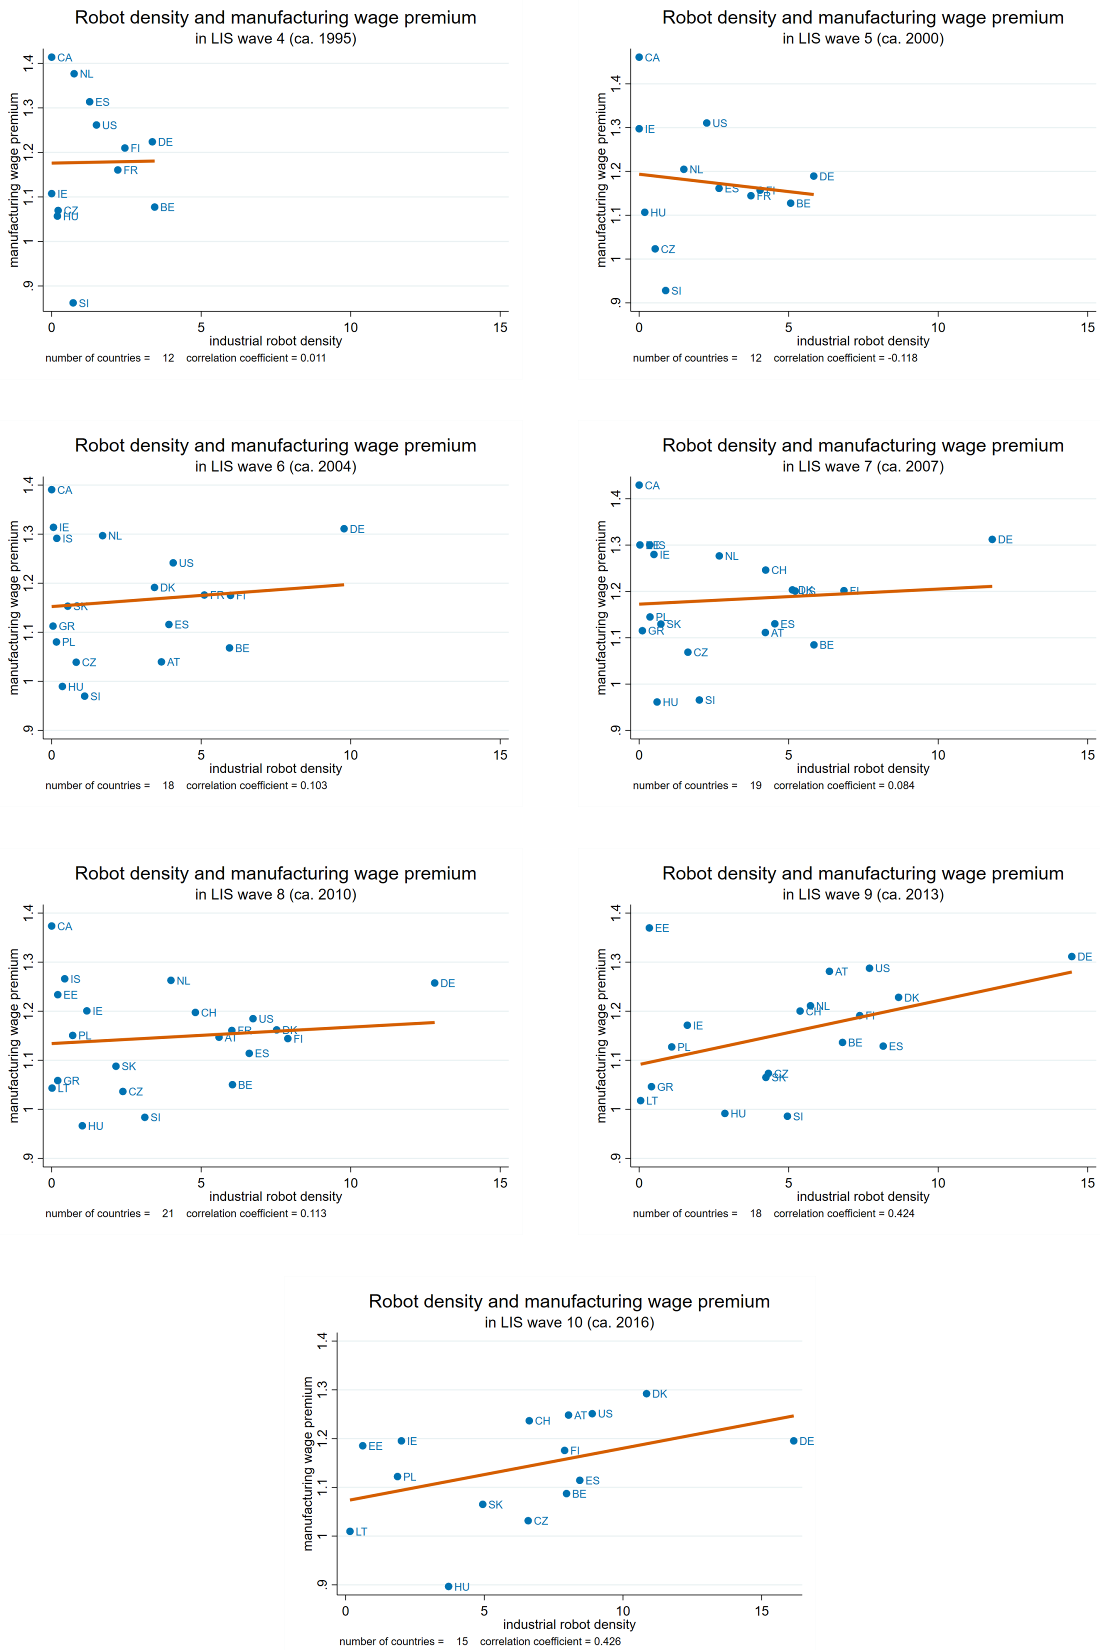

Figure D.2: Evolution of the relationship between robot density and the routine wage premium, by LIS wave. There is a trend towards a more positive correlation, although p-values do not reach conventional levels of statistical significance, partly due to low numbers.

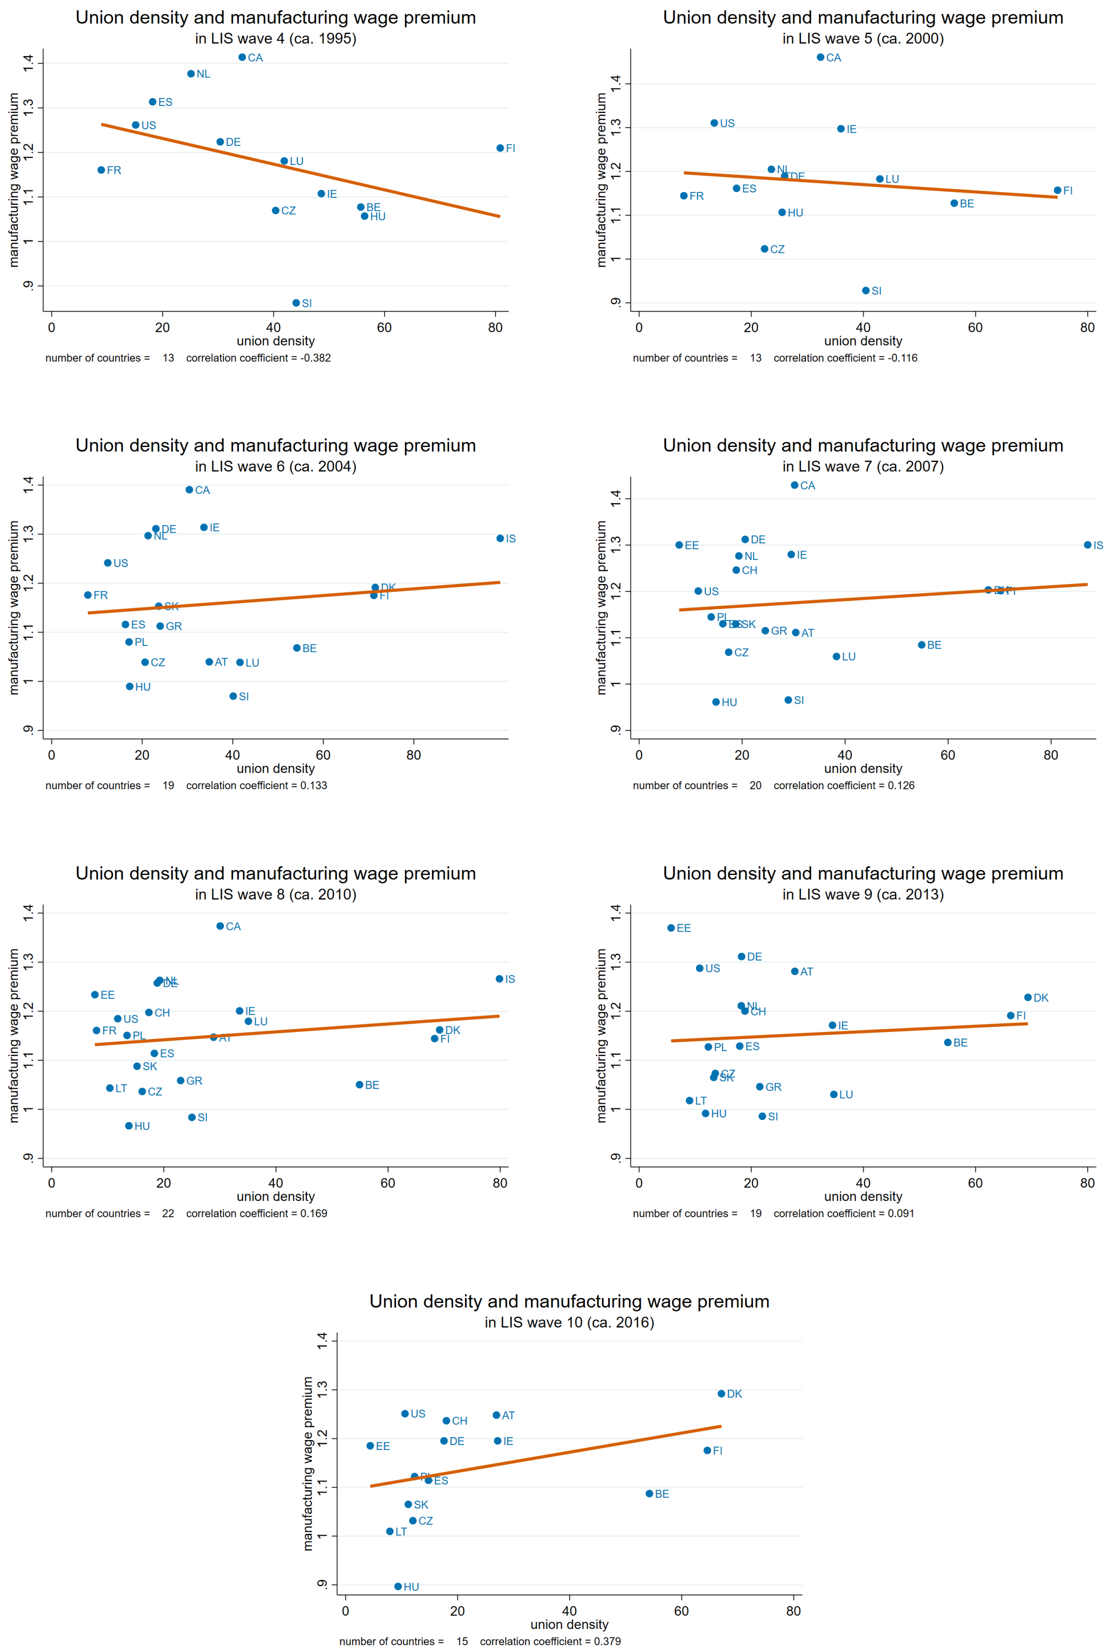

Figure D.3: Evolution of the relationship between union density and the routine wage premium, by LIS wave. There is a trend towards a more positive correlation, although p-values do not reach conventional levels of statistical significance.

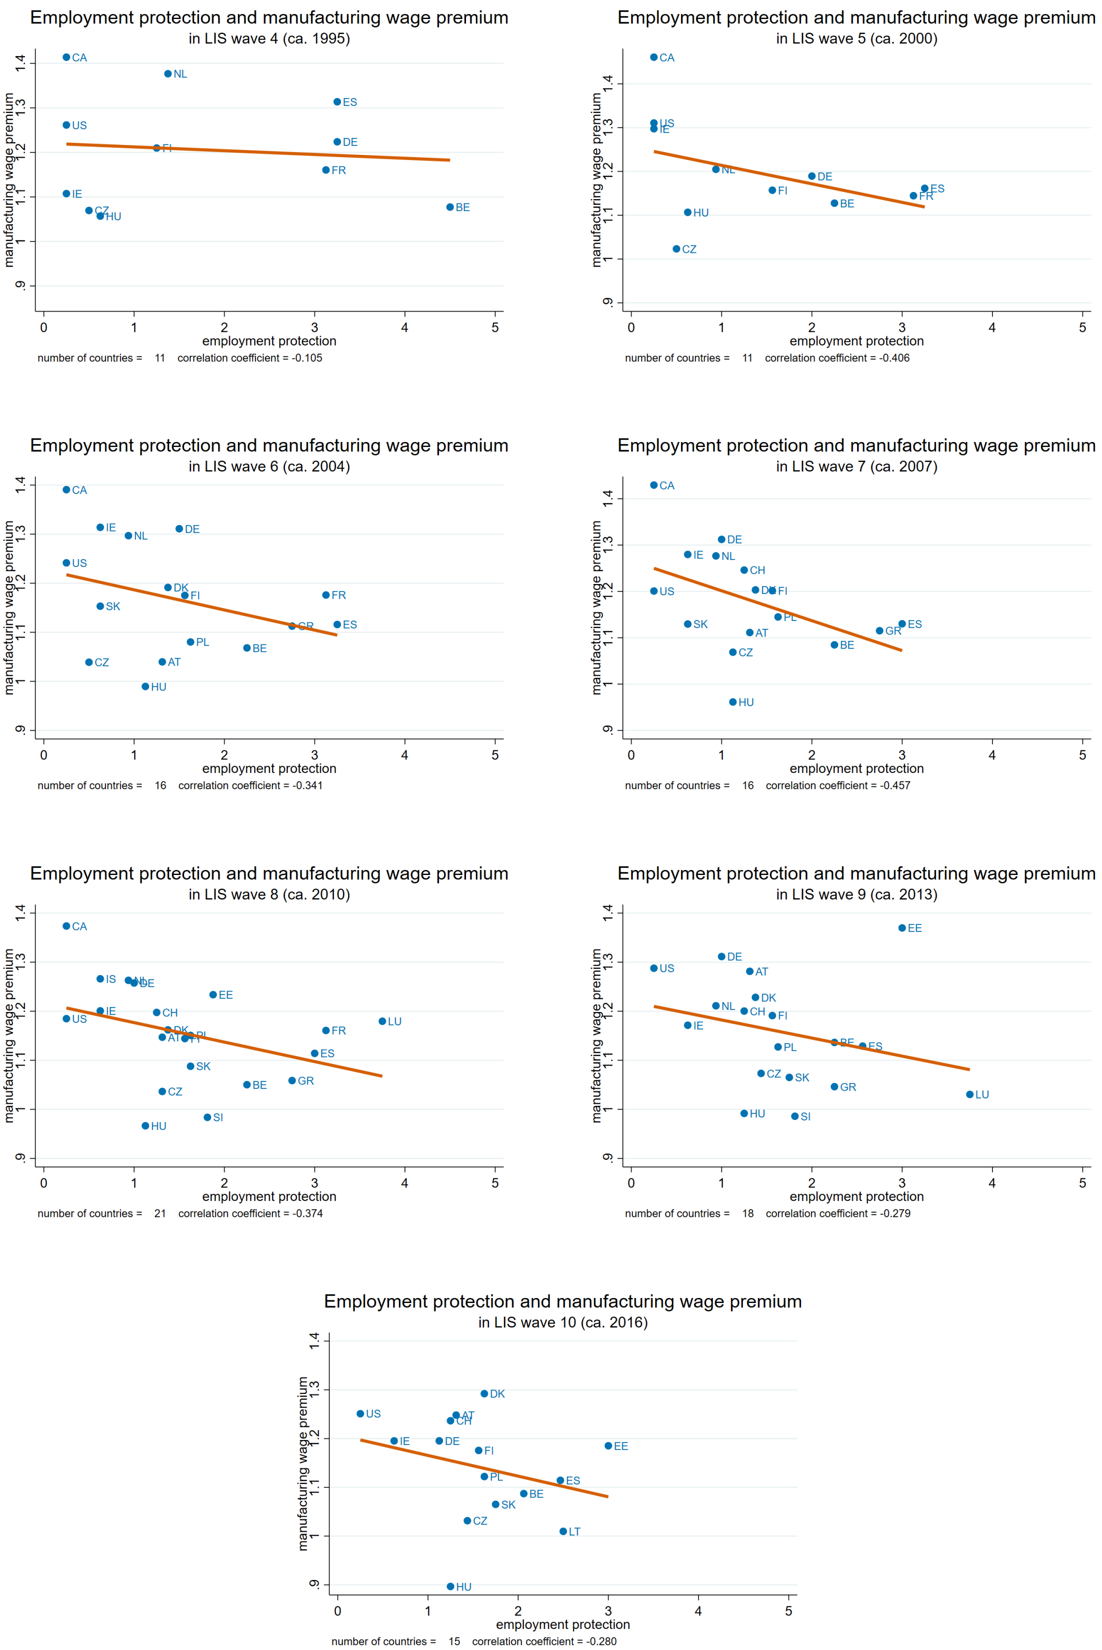

Figure D.4: Evolution of the relationship between employment protection and the routine wage premium, by LIS wave. There is a consistently negative correlation, although p-values do not reach conventional levels of statistical significance.

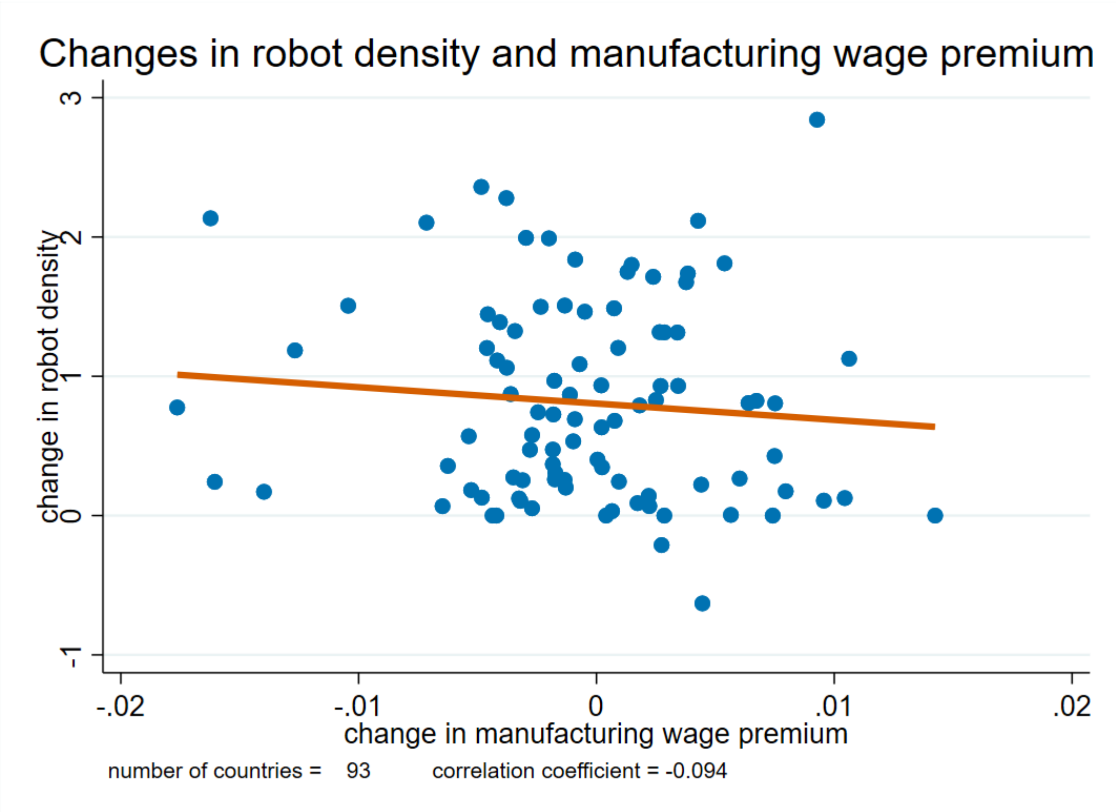

Figure D.5: *Wave-on-wave changes in robot density and the manufacturing wage premium. The negative correlation is not statistically significant.*

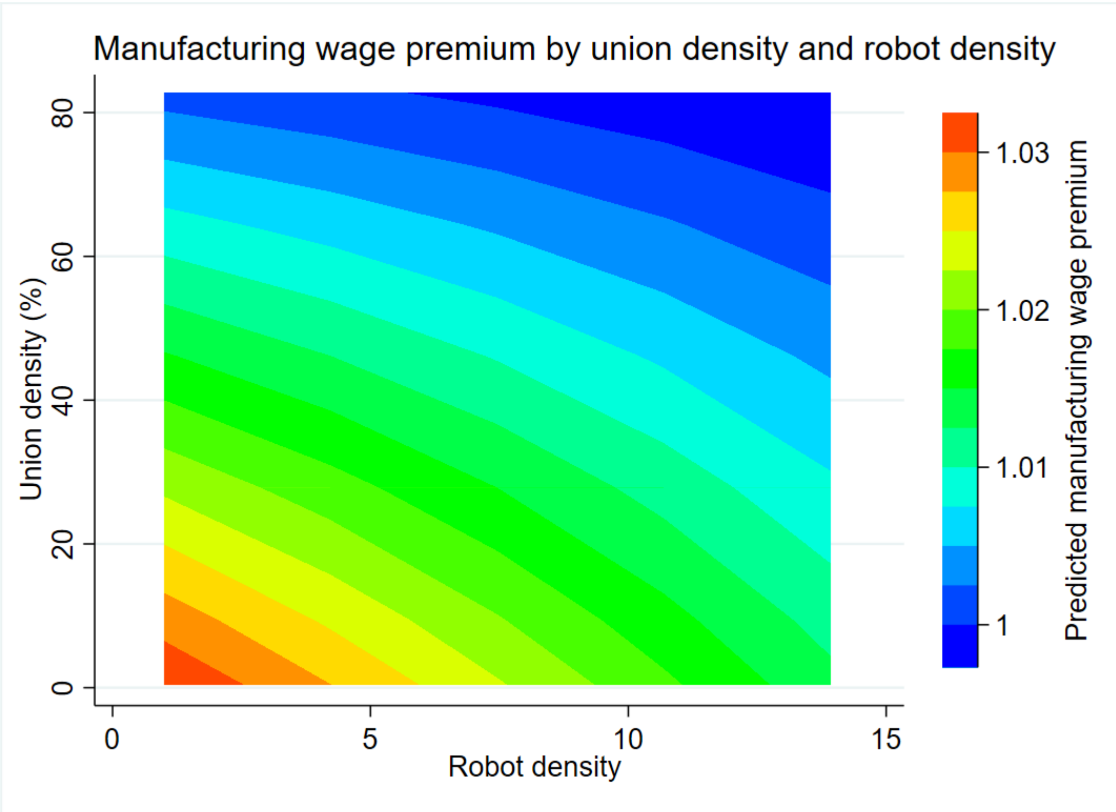

Figure D.6: *Predicted manufacturing wage premium by union density and robot density.*

Bibliography

Abadie, A., Athey, S., Imbens, G. & Wooldridge, J. (2017), When Should You Adjust Standard Errors for Clustering?, Technical report, National Bureau of Economic Research.
URL: *http://arxiv.org/abs/1710.02926*

Acemoglu, D. (1998), 'Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality', *The Quarterly Journal of Economics* **113**(4), 1055–1089.
URL: *https://doi.org/10.1162/003355398555838*

Acemoglu, D. (1999), 'Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence', *The American Economic Review* **89**(5), 1259–1278.

Acemoglu, D. (2002), 'Directed Technical Change', *The Review of Economic Studies* **69**(4), 781–809.

Acemoglu, D. (2003), 'Labor- and Capital-Augmenting Technical Change', *Journal of the European Economic Association* **1**(1), 1–37.

Acemoglu, D. & Autor, D. (2011), Skills, tasks and technologies: Implications for employment and earnings, *in* 'Handbook of Labor Economics', Vol. 4b, Elsevier B.V., pp. 1043–1171.

Acemoglu, D. & Autor, D. (2012), 'What Does Human Capital Do? A Review of Goldin and Katz's The Race between Education and Technology', *Journal of Economic Literature* **50**(2), 426–463.

Acemoglu, D., Autor, D., Dorn, D., Hanson, G. H. & Price, B. (2014), 'Return of the Solow Paradox? IT, Productivity, and Employment in US Manufacturing', *American Economic Review: Papers & Proceedings* **104**(5), 394–399.

Acemoglu, D., LeLarge, C. & Restrepo, P. (2020), Competing with Robots: Firm-Level Evidence from France.

Acemoglu, D. & Restrepo, P. (2018), 'Low-Skill and High-Skill Automation', *Journal of Human Capital* **12**(2), 204–232.

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

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

Acemoglu, D. & Restrepo, P. (2020*b*), 'The wrong kind of AI? Artificial intelligence and the future of labour demand', *Cambridge Journal of Regions, Economy and Society* **13**(1), 25–35.

Adermon, A. & Gustavsson, M. (2015), 'Job Polarization and Task-Biased Technological Change: Evidence from Sweden, 1975–2005', *Scandinavian Journal of Economics* **117**(3), 878–917.

Aghion, P., Antonin, C. & Bunel, S. (2019), 'Artificial intelligence, growth and employment: The role of policy', *Economie et Statistique* **2019**(510-512).

Aghion, P., Howitt, P. & Violante, G. L. (2002), 'General Purpose Technology and Wage Inequality', *Journal of Economic Growth* **7**(4), 315–345.

Ahlfeldt, G. M., Roth, D. & Seidel, T. (2018), 'The regional effects of Germany's national minimum wage', *Economics Letters* **172**, 127–130.
URL: *https://doi.org/10.1016/j.econlet.2018.08.032*

Andersen, J. G. (2012), Welfare States and Welfare State Theory.
URL: *http://vbn.aau.dk/ws/files/72613349/80_2012_J_rgen_Goul_Andersen.pdf*

Andrieu, E., Jamet, S., Marcolin, L. & Squicciarini, M. (2019), 'Occupational transitions: The cost of moving to a "safe haven"', *OECD Science, Technology and Innovation Policy Papers* **61**.

Anghel, B., De la Rica, S. & Lacuesta, A. (2014), 'The impact of the great recession on employment polarization in Spain', *SERIEs* **5**, 143–171.

Antonczyk, D., DeLeire, T. & Fitzenberger, B. (2010), 'Polarization and Rising Wage Inequality: Comparing the US and Germany', *IZA Discussion Paper Series* **4842**.
URL: *http://www.econstor.eu/handle/10419/30202*

Antonczyk, D., DeLeire, T. & Fitzenberger, B. (2018), 'Polarization and rising wage inequality: Comparing the U.S. and Germany', *Econometrics* **6**(20), 1–33.

Arestis, P., Ferreiro, J. & Gómez, C. (2020), 'Quality of employment and employment protection. Effects of employment protection on temporary and permanent employment', *Structural Change and Economic Dynamics* **53**, 180–188.
URL: *https://doi.org/10.1016/j.strueco.2020.02.008*

Armingeon, K., Engler, S. & Leemann, L. (2020), 'Comparative Political Data Set 1960 - 2018'.

Arntz, M., Gregory, T. & Zierahn, U. (2016), The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis.

Arntz, M., Gregory, T. & Zierahn, U. (2017), 'Revisiting the risk of automation', *Economics Letters* **159**, 157–160.
URL: *http://dx.doi.org/10.1016/j.econlet.2017.07.001*

Atkinson, A. B. (2015), *Inequality: What Can Be Done?*, Harvard University Press, Cambridge, MA.

Atkinson, A. B. & Stiglitz, J. E. (1969), 'A New View of Technological Change', *The Economic Journal* **79**(315), 573–578.

Autor, D. (2014), 'Skills, education, and the rise of earnings inequality among the "other 99 percent"', *Science* **344**(6186), 843–852.

Autor, D. (2015), 'Why Are There Still So Many Jobs? The History and Future of Workplace Automation', *Journal of Economic Perspectives* **29**(3), 3–30.
URL: *http://pubs.aeaweb.org/doi/10.1257/jep.29.3.3*

Autor, D. & 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.

Autor, D., Dorn, D. & Hanson, G. H. (2013), 'The China Syndrome: Local Labor Market Effects of Import Competition in the United States', *American Economic Reviewmerican Economic Review* **103**(6), 2121–2168.

Autor, D., Dorn, D. & Hanson, G. H. (2015), 'Untangling Trade and Technology: Evidence from Local Labour Markets', *The Economic Journal* **125**(May), 621–646.

Autor, D. H., Donohue III, J. J. & Schwab, S. J. (2006), 'The costs of wrongful-discharge laws', *The Review of Economics and Statistics* **88**(2), 211–231.

Autor, D. H., Kerr, W. R. & Kugler, A. D. (2007), 'Does Employment Protection Reduce Productivity? Evidence from US States', *The Economic Journal* **117**(June), 189–217.

Autor, D. & Handel, M. J. (2013), 'Putting Tasks to the Test: Human Capital, Job Tasks, and Wages', *Journal of Labor Economics* **31**(2), S59–S96.

Autor, D., Katz, L. F. & Kearney, M. S. (2008), 'Trends in U.S. Wage Inequality: Revising the Revisionists', *Review of Economics and Statistics* **90**(2), 300–323.
URL: *http://dx.doi.org/10.1162/rest.90.2.300*

Autor, D., Katz, L. F. & Krueger, A. B. (1998), 'Computing Inequality: Have Computers Changed the Labor Market?', *The Quarterly Journal of Economics* **113**(4), 1169–1213.

Autor, D., Levy, F. & Murnane, R. J. (2003), 'The Skill Content of Recent Technological Change: An Empirical Exploration', *The Quarterly Journal of Economics* **118**(4), 1279–1333.

Autor, D., Manning, A. & Smith, C. L. (2016), 'The contribution of the minimum wage to US wage inequality over three decades: A reassessment', *American Economic Journal: Applied Economics* **8**(1), 58–99.

Autor, D. & Salomons, A. (2018), 'Is Automation Labor Share-Displacing? Productivity Growth, Employment, and the Labor Share', *Brookings Papers on Economic Activity* **Spring 201**, 1–87.

Avent-Holt, D., Hällsten, M. & Cort, D. (2020), 'Occupational status and organizations: Variation in occupational hierarchies across Swedish workplaces', *Research in Social Stratification and Mobility* **70**(October 2018), 100423.
URL: *https://doi.org/10.1016/j.rssm.2019.100423*

Balsmeier, B. & Woerter, M. (2019), 'Is this time different? How digitalization influences job creation and destruction', *Research Policy* **48**(8), 103765.
URL: *https://doi.org/10.1016/j.respol.2019.03.010*

Bartel, A., Ichniowski, C. & Shaw, K. (2007), 'How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills', *The Quarterly Journal of Economics* **122**(4), 1721–1758.

Beaudry, P., Green, D. A. & Sand, B. M. (2014), 'The Declining Fortunes of the Young Since 2000', *American Economic Review: Papers & Proceedings* **104**(5), 381–386.

Beaudry, P., Green, D. A. & Sand, B. M. (2016), 'The Great Reversal in the Demand for Skill and Cognitive Tasks', *Journal of Labor Economics* **34**(1), S199–S247.
URL: *http://www.journals.uchicago.edu/doi/10.1086/682347*

Beck, N. (2001), 'Time-Series-Cross-Section Data: What Have We Learned in the Past Few Years ?', *Annual Review of Political Science* **4**, 271–93.

Beck, N. & Katz, J. N. (1995), 'What to Do (and Not to Do) with Time-Series Cross-Section Data', *The American Political Science Review* **100**(4), 634–647.

Bell, D. N. & Blanchflower, D. G. (2018), Underemployment in the US and Europe.
URL: *http://www.nber.org/papers/w24927*

Berman, E., Bound, J. & Griliches, Z. (1994), 'Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers', *The Quarterly Journal of Economics* **109**(2), 367–397.

Berman, E., Bound, J. & Machin, S. (1998), 'Implications of Skill-Biased Technological Change: International Evidence', *The Quarterly Journal of Economics* **113**(4), 1245–1279.

Bird, R. C. & Knopf, J. D. (2009), 'Do wrongful-discharge laws impair firm performance?', *Journal of Law and Economics* **52**(2), 197–222.

Blanchflower, D. G. & Bryson, A. (2004), 'What Effect Do Unions Have on Wages Now and Would Freeman and Medoff Be Surprised?', *Journal of Labor Research* **XXV**(3), 383–414.

Blau, F. D. & Kahn, L. M. (2002), *At home and abroad: U.S. labor-market performance in international perspective*, Russell Sage Foundation, New York.

Blau, F. D. & Kahn, L. M. (2017), 'The Gender Wage Gap: Extent, Trends, and Explanations', *Journal of Economic Literature* **55**(3), 789–865.

Blinder, A. S. (2009), 'How Many US Jobs Might Be Offshorable?', *World Economics* **10**(2), 41–78.

Blinder, A. S. & Krueger, A. B. (2013), 'Alternative measures of offshorability: A survey approach', *Journal of Labor Economics* **31**(2.2).

Boeri, T. (2011), Institutional Reforms and Dualism in European Labor Markets, *in* D. Card & O. Ashenfelter, eds, 'Handbook of Labor Economics', Vol. 4b, Elsevier B.V., chapter 13, pp. 1173–1236.

Boeri, T., Cahuc, P. & Zylberberg, A. (2015), 'The Costs of Flexibility- Enhancing Structural Reforms', *OECD Economics Department Working Papers* **1264**.

Boeri, T. & van Ours, J. (2008), *The Economics of Imperfect Labor Markets*, Princeton University Press, Princeton.

Bonfiglioli, A., Crinò, R., Gancia, G. & Papadakis, I. (2021), 'Robots, Offshoring and Welfare', *CEPR Discussion Paper Series* **16363**.

Bossler, M. & Gerner, H.-D. (2016), 'Employment effects of the new German minimum wage: Evidence from establishment-level micro data', *IAB Discussion Paper* **10**. **URL:** *http://doku.iab.de/discussionpapers/2016/dp1016.pdf*

Bostrom, N. (2014), *Superintelligence: paths, dangers, strategies*, Oxford University Press, Oxford.

Bound, J. & Johnson, G. (1992), 'Changes in the Structure of Wages in the 1980's : An Evaluation of Alternative Explanations', *The American Economic Review* **82**(3), 371–392.

Brady, D., Baker, R. S. & Finnigan, R. (2013), 'When Unionization Disappears: State-Level Unionization and Working Poverty in the United States', *American Sociological Review* **78**(5), 872–896.

Breda, T. (2015), 'Firms' Rents, Workers' Bargaining Power and the Union Wage Premium', *The Economic Journal* **125**(December), 1616–1652.

Bresnahan, T. F., Brynjolfsson, E. & Hitt, L. M. (2002), 'Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence', *The Quarterly Journal of Economics* **117**(1), 339–376.

Bresnahan, T. F. & Trajtenberg, M. (1995), 'General purpose technologies: 'Engines of growth'?', *Journal of Econometrics* **65**(1), 83–108.

Brown, C. (1999), Minimum Wages, Employment, and the Distribution of Income, *in* O. Ashenfelter & D. Card, eds, 'Handbook of Labor Economics', Vol. 3b, Elsevier B.V., chapter 32, pp. 2101–2163.

Brynjolfsson, E. & McAfee, A. (2014), *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*, W. W. Norton & Co., New York.

Bryson, A., Ebbinghaus, B. & Visser, J. (2011), 'Introduction: Causes, consequences and cures of union decline', *European Journal of Industrial Relations* **17**(2), 97–105.

Bukodi, E. & Goldthorpe, J. H. (2019), *Social Mobility and Education in Britain: Research, Politics and Policy*, Cambridge University Press, Cambridge.

Bureau of Labor Statistics (2021), Union Members - 2020, Technical report, Bureau of Labor Statistics.
URL: *https://www.bls.gov/news.release/pdf/union2.pdf*

Burgoon, B. (2001), 'Globalization and Welfare Compensation: Disentangling the Ties that Bind', *International Organization* **55**(3), 509–551.

Caines, C., Hoffmann, F. & Kambourov, G. (2017), 'Complex-task biased technological change and the labor market', *Review of Economic Dynamics* **25**, 298–319.
URL: *http://dx.doi.org/10.1016/j.red.2017.01.008*

Card, D. & DiNardo, J. (2002), 'Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles', *Journal of Labor Economics* **20**(4), 733–783.
URL: *http://www.journals.uchicago.edu/doi/10.1086/342055*

Card, D. & Krueger, A. B. (1994), 'Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania', *The American Economic Review* **84**(4), 772–793.

Card, D., Lemieux, T. & Riddell, W. C. (2003), Unionization and the Wage Structure, *in* J. T. Addison & C. Schnabel, eds, 'International Handbook of Trade Unions', Edward Elgar, Cheltenham, pp. 246–292.

Card, D., Lemieux, T. & Riddell, W. C. (2004), 'Unions and Wage Inequality', *Journal of Labor Research* **XXV**(4), 519–562.

Cardullo, G., Conti, M. & Sulis, G. (2015), 'Sunk capital, unions and the hold-up problem: Theory and evidence from cross-country sectoral data', *European Economic Review* **76**, 253–274.
URL: *http://dx.doi.org/10.1016/j.euroecorev.2015.03.005*

Cengiz, D., Dube, A., Lindner, A. & Zipperer, B. (2019), 'The Effect of Minimum Wages on Low-Wage Jobs', *The Quarterly Journal of Economics* pp. 1405–1454.

Cerina, F., Moro, A. & Rendall, M. (2021), A Note on Employment and Wage Polarization in the U.S.

Checchi, D. & García-Penalosa, C. (2008), 'Labour Market Institutions and Income Inequality', *Economic Policy* **23**(56), 601–649.

Chiacchio, F., Petropoulos, G. & Pichler, D. (2018), 'The Impact of Industrial Robots on EU Employment and Wages: A Local Labour Market Approach', *Bruegel* pp. 1–28.
URL: *https://www.bruegel.org/2018/04/the-impact-of-industrial-robots-on-eu-employment-and-wages-a-local-labour-market-approach/*

CIA (2021), 'The World Factbook'.
URL: *https://www.cia.gov/the-world-factbook/countries/*

Ciminelli, G., Duval, R. & Furceri, D. (2018), 'Employment Protection Deregulation and Labor Shares in Advanced Economies', *IMF Working Papers* **18**(186).

Cirillo, V. (2018), 'Job polarization in European industries', *International Labour Review* **157**(1), 39–63.

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.
URL: *http://www.journals.uchicago.edu/doi/10.1086/682289*

Cortes, G. M., Jaimovich, N. & Siu, H. E. (2017), 'Disappearing routine jobs: Who, how, and why?', *Journal of Monetary Economics* **91**, 69–87.

Cortes, G. M., Jaimovich, N. & Siu, H. E. (2018), The "End of Men" and Rise of Women in the High-Skill Labor Market.

Cortes, G. M., Nekarda, C. J., Jaimovich, N. & Siu, H. E. (2020), 'The dynamics of disappearing routine jobs: A flows approach', *Labour Economics* **65**(March), 101823.
URL: *https://doi.org/10.1016/j.labeco.2020.101823*

Cortes, G. M. & Salvatori, A. (2019), 'Delving into the demand side: Changes in workplace specialization and job polarization', *Labour Economics* **57**(February), 164–176.
URL: *https://doi.org/10.1016/j.labeco.2019.02.004*

Cséfalvay, Z. (2020), 'Robotization in Central and Eastern Europe: catching up or dependence?', *European Planning Studies* **28**(8), 1534–1553.

Damiani, M., Pompei, F. & Ricci, A. (2020), 'Labour shares, employment protection and unions in European economies', *Socio-Economic Review* **18**(4), 1001–1038.

Dauth, W., Findeisen, S., Suedekum, J. & Woessner, N. (2019), The Adjustment of Labor Markets to Robots.

De Boef, S. & Keele, L. (2008), 'Taking Time Seriously', *American Journal of Political Science* **52**(1), 184–200.

Denice, P. & Rosenfeld, J. (2018), 'Unions and Nonunion Pay in the United States, 1977-2015', *Sociological Science* **5**, 541–561.

Dickens, R., Machin, S. & Manning, A. (1999), 'The Effects of Minimum Wages on Employment: Theory and Evidence from Britain', *Journal of Labor Economics* **17**(1), 1–22.

Dickens, R. & Manning, A. (2004), 'Has the National Minimum Wage Reduced UK Wage Inequality?', *Journal of the Royal Statistical Society. Series A (Statistics in Society)* **167**(4), 613–626.

DiNardo, J., Fortin, N. M. & Lemieux, T. (1996), 'Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach', *Econometrica* **64**(5), 1001–1044.

DiPrete, T. A., De Graaf, P. M., Luijkx, R., Tåhlin, M. & Blossfeld, H. P. (1997), 'Collectivist versus individualist mobility regimes? Structural change and job mobility in four countries', *American Journal of Sociology* **103**(2), 318–358.

Dolton, P. & Marcenaro-Gutierrez, O. D. (2011), 'If you pay peanuts do you get monkeys? A cross-country analysis of teacher pay and pupil performance', *Economic Policy* **January**, 5–55.

Dribbusch, H., Lehndorff, S. & Schulten, T. (2018), Two worlds of unionism? German manufacturing and service unions since the Great Recession, *in* S. Lehndorff, H. Dribbusch & T. Schulten, eds, 'Rough waters: European trade unions in a time of crises', European Trade Union Institute, Brussels, pp. 197–220.

Dube, A., Lester, T. W. & Reich, M. (2010), 'Minimum wage effects across state borders: Estimates using contiguous counties', *Review of Economics and Statistics* **92**(4), 945–964.

Dube, A. & Lindner, A. (2021), 'City limits: What do local-area minimum wages do?', *Journal of Economic Perspectives* **35**(1), 27–50.

Dustmann, C., Fitzenberger, B., Schönberg, U. & Spitz-Oener, A. (2014), 'From Sick Man of Europe to Economic Superstar: Germany's Resurgent Economy', *Journal of Economic Perspectives* **28**(1), 167–188.
URL: *http://pubs.aeaweb.org/doi/10.1257/jep.28.1.167*

Dustmann, C., Lindner, A., Schönberg, U., Umkehrer, M. & Berge, P. (2021), 'Reallocation Effects of the Minimum Wage', *The Quarterly Journal of Economics* .
URL: *https://www.cream-migration.org/publ_uploads/CDP_07_20.pdf*

Dustmann, C., Ludsteck, J. & Schönberg, U. (2009), 'Revisiting the German Wage Structure', *The Quarterly Journal of Economics* **124**(2), 843–881.

Dwyer, R. E. (2013), 'The Care Economy? Gender, Economic Restructuring, and Job Polarization in the U.S. Labor Market', *American Sociological Review* **78**(3), 390–416.

Dwyer & Wright (2019), 'Low-Wage Job Growth, Polarization, and the Limits and Opportunities of the Service Economy', *RSF: The Russell Sage Foundation Journal of the Social Sciences* **5**(4), 56.

Ebbinghaus, B. & Visser, J. (1999), 'When Institutions Matter: Union Growth and Decline in Western Europe, 1950-95', *European Sociological Review* **15**(2), 135–158.

Ebbinghaus, B. & Visser, J. (2000), *Trade Unions in Western Europe since 1945*, Macmillan, Basingstoke.

Eichhorst, W. & Marx, P. (2015), Introduction: An occupational perspective on non-standard employment, *in* W. Eichhorst & P. Marx, eds, 'Non-Standard Employment in Post-Industrial Labour Markets: An Occupational Perspective', Edward Elgar Publishing, Cheltenham, pp. 1–25.

Elias, P. (1997), 'Occupational classification (ISCO-88): Concepts, methods, reliability, validity and cross-national comparability', *OECD Labour Market and Social Policy Occasional Papers* **20**, 1–22.

Elliott, S. W. (2017), Projecting the Impact of Information Technology on Work and Skills in the 2030s, *in* J. Buchanan, D. Finegold, K. Mayhew & C. Warhurst, eds, 'The Oxford Handbook of Skills and Training', Oxford University Press, Oxford.
URL:
http://oxfordhandbooks.com/view/10.1093/oxfordhb/9780199655366.001.0001/oxfordhb-9780199655366-e-26

Esping-Andersen, G. (1990), *The Three Worlds of Welfare Capitalism*, Polity Press, Cambridge.

Eurofound (2014), *Drivers of recent job polarisation and upgrading in Europe: European Jobs Monitor 2014*, Publications Office of the European Union, Luxembourg.

Eurofound (2017), Occupational change and wage inequality: European Jobs Monitor 2017, Technical report, Publications Office of the European Union, Luxembourg.
URL: *https://www.eurofound.europa.eu/publications/report/2017/occupational-change-and-wage-inequality-european-jobs-monitor-2017*

Eurofound (2021), 'European Working Conditions Survey'.
URL:
https://www.eurofound.europa.eu/surveys/european-working-conditions-surveys-ewcs

Feldmann, H. (2008), 'Business regulation and labor market performance around the world', *Journal of Regulatory Economics* **33**(2), 201–235.

Feldmann, H. (2009), 'The unemployment effects of labor regulation around the world', *Journal of Comparative Economics* **37**(1), 76–90.
URL: *http://dx.doi.org/10.1016/j.jce.2008.10.001*

Fernández-Macías, E. (2012), 'Job Polarization in Europe? Changes in the Employment Structure and Job Quality, 1995-2007', *Work and Occupations* **39**(2), 157–182.

Fernández-Macías, E. & Hurley, J. (2017), 'Routine-biased technical change and job polarization in Europe', *Socio-Economic Review* **15**(3), 563–585.

Fernández, R. (2013), 'Cultural change as learning: The evolution of female labor force participation over a century', *American Economic Review* **103**(1), 472–500.

Fernandez, R. M. (2001), 'Skill-biased technological change and wage inequality: Evidence from a plant retooling', *American Journal of Sociology* **10**(2), 273–320.

Firpo, S., Fortin, N. M. & Lemieux, T. (2011), Occupational Tasks and Changes in the Wage Structure.

Fishback, P. V. & Kantor, S. E. (1995), 'Did Workers Pay for the Passage of Workers' Compensation Laws?', *The Quarterly Journal of Economics* **110**(3), 713–742.

Fonseca, T., Lima, F. & Pereira, S. C. (2018), 'Job polarization, technological change and routinization: Evidence for Portugal', *Labour Economics* **51**(February), 317–339.
URL: *https://doi.org/10.1016/j.labeco.2018.02.003*

Förster, M. F. & Tóth, I. G. (2015), Cross-Country evidence of the multiple causes of inequality changes in the OECD area, *in* A. B. Atkinson & F. Bourguignon, eds, 'Handbook of Income Distribution', Vol. 2B, Elsevier B.V., pp. 1729–1843.

Freeman, R. B. (2008), Why Do We Work More Than Keynes Expected?, *in* L. Pecchi & G. Piga, eds, 'Revisiting Keynes: Economic Possibilities for Our Grandchildren', MIT Press Scholarship Online.

Freeman, R. B. & Medoff, J. L. (1984), *What Do Unions Do?*, Basic Books, New York.

Frey, C. B. & Osborne, M. A. (2017), 'The future of employment: How susceptible are jobs to computerisation?', *Technological Forecasting & Social Change* **114**, 254–280.
URL: *http://dx.doi.org/10.1016/j.techfore.2016.08.019*

Frey, C. B., Osborne, M. A. & Holmes, C. (2016), Technology at Work v2.0: The Future Is Not What It Used to Be, Technical report, Citigroup, London.

Gallie, D. (1991), 'Patterns of Skill Change - Upskilling, Deskilling or the Polarization of Skills', *Work, Employment and Society* **5**(3), 319–351.

Gallie, D. (2017), 'The Quality of Work in a Changing Labour Market', *Social Policy and Administration* **51**(2), 226–243.

Gallie, D., Felstead, A. & Green, F. (2012), 'Job preferences and the intrinsic quality of work: The changing attitudes of British employees 1992-2006', *Work, Employment and Society* **26**(5), 806–821.

Ganzeboom, H. B. & Treiman, D. J. (1996), 'Internationally comparable measures of occupational status for the 1988 International Standard Classification of Occupations', *Social Science Research* **25**(3), 201–239.

Goldfarb, A. & Tucker, C. (2019), 'Digital economics', *Journal of Economic Literature* **57**(1), 3–43.

Goldin, C. (2014), 'A grand gender convergence: Its last chapter', *American Economic Review* **104**(4), 1091–1119.

Goldin, C. & Katz, L. F. (2008), *The Race between Education and Technology*, Harvard University Press, Cambridge, MA.

Goldthorpe, J. H. (2013), 'Understanding - And misunderstanding - Social mobility in Britain: The entry of the economists, the confusion of politicians and the limits of educational policy', *Journal of Social Policy* **42**(3), 431–450.

Goldthorpe, J. H., Llewellyn, C. & Payne, C. (1987), *Social mobility and class structure in modern Britain*, 2nd ed. edn, Clarendon Press, Oxford.

Goldthorpe, J. H. & McKnight, A. (2004), The Economic Basis of Social Class.

Goos, M. & Manning, A. (2007), 'Lousy and Lovely Jobs : The Rising Polarization of Work in Britain', *The Review of Economics and Statistics* **89**(1), 118–133.

Goos, M., Manning, A. & Salomons, A. (2009), 'Job Polarization in Europe', *American Economic Review: Papers & Proceedings* **99**(2), 59–63.

Goos, M., Manning, A. & Salomons, A. (2014), 'Explaining Job Polarization: Routine-Biased Technological Change and Offshoring', *American Economic Review* **104**(8), 2509–2526.
URL: *http://pubs.aeaweb.org/doi/10.1257/aer.104.8.2509*

Graetz, G. & Michaels, G. (2018), 'Robots at Work', *Review of Economics and Statistics* **C**(5), 753–768.

Grossman, G. M. & Rossi-Hansberg, E. (2008), 'Trading tasks: A simple theory of offshoring', *American Economic Review* **98**(5), 1978–1997.

Handel, M. J. (2017), Measuring Job Content: Skills, Technology, and Management Practices, *in* J. Buchanan, D. Finegold, K. Mayhew & C. Warhurst, eds, 'The Oxford Handbook of Skills and Training', Oxford University Press, Oxford.
URL: *http://oxfordhandbooks.com/view/10.1093/oxfordhb/9780199655366.001.0001/oxfordhb-9780199655366-e-5*

Haslberger, M. (2021*a*), 'Rethinking the measurement of occupational task content', *Economic and Labour Relations Review* pp. 1–22.
URL: *https://doi.org/10.1177/10353046211037095*

Haslberger, M. (2021*b*), 'Routine-biased technological change does not always lead to polarisation: Evidence from 10 OECD countries, 1995–2013', *Research in Social Stratification and Mobility* **74**(August).
URL: *https://linkinghub.elsevier.com/retrieve/pii/S0276562421000433*

Hershbein, B. & Kahn, L. B. (2018), 'Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings', *American Economic Review* **108**(7), 1737–1772.

Heyes, J. & Tomlinson, M. (2021), 'Underemployment and well-being in Europe', *Human Relations* **74**(8), 1240–1266.

Heyman, F. (2016), 'Job polarization , job tasks and the role of firms', *Economics Letters* **145**, 246–251.
URL: *http://dx.doi.org/10.1016/j.econlet.2016.06.032*

Holmes, C. (2017), The Individual Benefits of Investing in Skills, *in* J. Buchanan, D. Finegold, K. Mayhew & C. Warhurst, eds, 'The Oxford Handbook of Skills and Training', Oxford University Press, Oxford.
URL: *http://oxfordhandbooks.com/view/10.1093/oxfordhb/9780199655366.001.0001/oxfordhb-9780199655366-e-17*

Horemans, J., Marx, I. & Nolan, B. (2016), 'Hanging in, but only just: part-time employment and in-work poverty throughout the crisis', *IZA Journal of European Labor Studies* **5**(5).
URL: *http://dx.doi.org/10.1186/s40174-016-0053-6*

Hout, M. & DiPrete, T. A. (2006), 'What we have learned: RC28's contributions to knowledge about social stratification', *Research in Social Stratification and Mobility* **24**(1), 1–20.

Hunt, J. & Nunn, R. (2019), 'Is employment polarization informative about wage inequality and is employment really polarizing?', *NBER Working Paper Series* (26064).

Hutter, C. & Weber, E. (2017), 'Labour Market Effects of Wage Inequality and Skill Biased Technical Change in Germany', *IAB Discussion Paper* **5**.

International Organization for Standardization (2012), 'ISO 8373:2012(en): Robots and Robotic Devices - Vocabulary'.
URL: *https://www.iso.org/obp/ui/#iso:std:iso:8373:ed-2:v1:en*

Jardim, E., Long, M., Plotnick, R., van Inwegen, E., Vigdor, J. & Wething, H. (2017), 'Minimum Wage Increases, Wages, and Low-Wage Employment: Evidence from Seattle', *NBER Working Paper Series* **23532**.
URL: *http://www.nber.org/papers/w23532.pdf*

Jerbashian, V. (2019), 'Automation and Job Polarization: On the Decline of Middling Occupations in Europe', *Oxford Bulletin of Economics and Statistics* .

Kahn, L. M. (2010), 'Employment protection reforms, employment and the incidence of temporary jobs in Europe: 1996-2001', *Labour Economics* **17**(1), 1–15.
URL: *http://dx.doi.org/10.1016/j.labeco.2009.05.001*

Katz, L. F. & Autor, D. (1999), Changes in the Wage Structure and Earnings Inequality, *in* D. Card & O. Ashenfelter, eds, 'Handbook of Labour Economics', Vol. 3a, Elsevier Science B.V., chapter 26, pp. 1463–1555.

Katz, L. F. & Murphy, K. M. (1992), 'Changes in Relative Wages, 1963-1987: Supply and Demand Factors', *The Quarterly Journal of Economics* **107**(1), 35–78.

Keele, L. & Kelly, N. J. (2006), 'Dynamic Models for Dynamic Theories: The Ins and Outs of Lagged Dependent Variables', *Political Analysis* **14**, 186–205.

Keynes, J. M. (2010), Economic Possibilities for Our Grandchildren, *in* 'Essays in Persuasion', Palgrave Macmillan, London.

Klenert, D., Fernández-Macías, E. & Antón, J.-I. (2020), Do robots really destroy jobs? Evidence from Europe, Technical report, Joint Research Centre, Sevilla.

Koeniger, W., Leonardi, M. & Nunziata, L. (2007), 'Labor Market Institutions and Wage Inequality', *Industrial and Labor Relations Review* **60**(3), 340–356.

Kollmeyer, C. (2009), 'Explaining deindustrialization: How affluence, productivity growth, and globalization diminish manufacturing employment', *American Journal of Sociology* **114**(6), 1644–1674.

Korpi, W. (1983), *The Democratic Class Struggle*, Routledge & Kegan Paul, London.

Korpi, W. (1985), 'Power Resources Approach vs. Action and Conflict: On Causal and Intentional Explanations in the Study of Power', *Sociological Theory* **3**(2), 31.

Korpi, W. & Palme, J. (1998), 'The Paradox of Redistribution and Strategies of Equality: Welfare State Institutions, Inequality, and Poverty in the Western Countries', *American Sociological Review* **63**(5), 661–687.

Kristal, T. & Cohen, Y. (2015), 'What do computers really do? Computerization, fading pay-setting institutions and rising wage inequality', *Research in Social Stratification and Mobility* **42**, 33–47.
URL: *http://dx.doi.org/10.1016/j.rssm.2015.07.001*

Kristal, T. & Cohen, Y. (2017), 'The causes of rising wage inequality: The race between institutions and technology', *Socio-Economic Review* **15**(1), 187–212.

Krueger, A. (1991), 'How computers have changed the wage strucutre: evidence from microdata, 1984-89', *NBER Working Paper Series* **3858**.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V. & Violante, G. L. (2000), 'Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis', *Econometrica* **68**(5), 1029–1053.

le Grand, C. & Tåhlin, M. (2013), 'Class, occupation, wages, and skills: The iron law of labor market inequality', *Comparative Social Research* **30**, 3–46.

Lee, D. S. (1999), 'Wage Inequality in the United States during the 1980s: Rising Dispersion or Falling Minimum Wage?', *The Quarterly Journal of Economics* **114**(3), 977–1023.

Lehndorff, S., Dribbusch, H. & Schulten, T. (2018), European trade unions in a time of crises - an overview, *in* S. Lehndorff, H. Dribbusch & T. Schulten, eds, 'Rough waters: European trade unions in a time of crises', ETUI Printshop, Brussels, pp. 7–38.

Lemieux, T. (2006), 'Increasing Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill?', *The American Economic Review* **96**(3), 461–498.

Lemieux, T. (2008), 'The Changing Nature of Wage Inequality', *Journal of Population Economics* **21**(1), 21–48.

Leonard, J. S. (1992), 'Unions and Employment Growth', *Industrial Relations* **31**(1).

Levy, F. & Murnane, R. J. (1992), 'U.S. earnings levels and earnings inequality: a review of recent trends and proposed explanations', *Journal of Economic Literature* **30**(3), 1333–1381.
URL: *http://www.jstor.org/stable/2728062*

LIS (2019), 'Generic Codebook: LIS Database (2019 Template)'.
URL: *https://www.lisdatacenter.org/frontend#/home*

Liu, Y. & Grusky, D. B. (2013), 'The payoff to skill in the third industrial revolution', *American Journal of Sociology* **118**(5), 1330–1374.

Longmuir, M., Schröder, C. & Targa, M. (2020), 'De-routinization of jobs and polarization of earnings: Evidence from 35 countries', *LIS Working Paper Series* (796).

Lordan, G. & Neumark, D. (2018), 'People versus machines: The impact of minimum wages on automatable jobs', *Labour Economics* **52**(March), 40–53.
URL: *http://linkinghub.elsevier.com/retrieve/pii/S0927537118300228*

Luxembourg Income Study (2021), 'LIS Database'.
URL: *https://www.lisdatacenter.org/our-data/lis-database/*

MacCrory, F., Westerman, G., Alhammadi, Y. & Brynjolfsson, E. (2016), Racing With and Against the Machine: Changes in Occupational Skill Composition in an Era of Rapid Technological Advance, *in* 'Thirty Fifth International Conference on Information Systems, Auckland'.

Machin, S. (2016), 'Rising wage inequality, real wage stagnation and unions', *Research in Labor Economics* **43**, 329–354.

Machin, S., Manning, A. & Rahman, L. (2003), 'Where the Minimum Wage Bites Hard: Introduction of Minimum Wages to a Low Wage Sector', *Journal of the European Economic Association* **1**(1), 154–180.

MacLeod, W. B. (2011), Great Expectations: Law, Employment Contracts, and Labor Market Performance, *in* D. Card & O. Ashenfelter, eds, 'Handbook of Labor Economics', Vol. 4b, Elsevier B.V., chapter 18, pp. 1591–1696.

Mahutga, M. C., Curran, M. & Roberts, A. (2018), 'Job tasks and the comparative structure of income and employment: Routine task intensity and offshorability for the LIS', *International Journal of Comparative Sociology* **59**(2), 81–109.

Manning, A. (2004), 'We Can Work It Out: the Impact of Technological Change on the Demand for Low Skill Workers', *Scottish Journal of Political Economy* **51**(5), 581–609.

Mazzolari, F. & Ragusa, G. (2013), 'Spillovers from high-skill consumption to low-skill labor markets', *The Review of Economics and Statstics* **95**(1), 74–86.

Michaels, G., Natraj, A. & Van Reenen, J. (2014), 'Has ICT Polarized Skill Demand? Evidence From Eleven Countries Over Twenty-Five Years', *The Review of Economics and Statstics* **96**(March), 60–77.
URL: *http://www.mitpressjournals.org/doi/pdf/10.1162/REST_a_00366*

Mincer, J. (1991), 'Human Capital, Technology, and the Wage Structure: What Do Time Series Show?', *NBER Working Paper Series* **3581**.

Mishel, L. & Bivens, J. (2017), The zombie robot argument lurches on - There is no evidence that automation leads to joblessness or inequality, Technical Report May, Economic Policy Institute.
URL: *http://www.epi.org/files/pdf/126750.pdf*

Mishel, L., Bivens, J., Gould, E. & Shierholz, H. (2013), *The State of Working America*, 12 edn, Cornell University Press, Ithaca.

Mishel, L., Schmitt, J. & Shierholz, H. (2014), 'Wage Inequality: A Story of Policy Choices', *New Labor Forum* **23**(3), 26–31.
URL: *http://journals.sagepub.com/doi/10.1177/1095796014544325*

Mishel, L., Shierholz, H. & Schmitt, J. (2013), Don't Blame the Robots: Assessing the Job Polarization Explanation of Growing Wage Inequality.

Mokyr, J., Vickers, C. & Ziebarth, N. L. (2015), 'The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different?', *Journal of Economic Perspectives* **29**(3), 31–50.
URL: *http://pubs.aeaweb.org/doi/10.1257/jep.29.3.31*

Moretti, E. (2012), *The New Geography of Jobs*, Mariner Books, New York.

Mouw, T. & Kalleberg, A. L. (2010), 'Occupations and the structure of wage inequality in the united states, 1980s to 2000s', *American Sociological Review* **75**(3), 402–431.

Murphy, E. C. & Oesch, D. (2018), 'Is Employment Polarisation Inevitable? Occupational Change in Ireland and Switzerland, 1970-2010', *Work, Employment and Society* **32**(6), 1099–1117.
URL: *https://doi.org/10.1177/0950017017738944*

Murphy, K. M. & Topel, R. H. (2016), 'Human Capital Investment, Inequality and Economic Growth', *Journal of Labor Economics* **34**(2), S99–S127.
URL: *http://www.nber.org/papers/w21841*

Naticchioni, P., Ragusa, G. & Massari, R. (2014), 'Unconditional and Conditional Wage Polarization in Europe', *IZA Discussion Paper* **8465**.

Neumark, D., Schweitzer, M. & Wascher, W. (2004), 'Minimum Wage Effects throughout the Wage Distribution', *The Journal of Human Resources* **39**(2), 425–450.

Nickell, S. & Andrews, M. (1983), 'Unions, Real Wages and Employment in Britain 1951-79', *Oxford Economic Papers* **35**(Supplement: The Causes of Unemployment), 183–206.

Nolan, B., Rahbari, E., Richiardi, M., Valenzuela Riviera, L. & Nabarro, B. (2017), Inequality and Prosperity in the Industrialized World: Addressing a Growing Challenge, Technical report, Citigroup.

Nolan, B. & Valenzuela, L. (2019), 'Inequality and its discontents', *Oxford Review of Economic Policy* **35**(3), 396–430.

Nolan, B. & Voitchovsky, S. (2016), 'Job loss by wage level: lessons from the Great Recession in Ireland', *IZA Journal of European Labor Studies* **5**(7).
URL: *http://dx.doi.org/10.1186/s40174-016-0057-2*

Nordhaus, W. D. (2007), 'Two Centuries of Productivity Growth in Computing', *The Journal of Economic History* **67**(01), 128–159.
URL: *http://www.journals.cambridge.org/abstract_S0022050707000058*

Nunziata, L. (2005), 'Institutions and Wage Determination: a Multi-country Approach', *Oxford Bulletin of Economics and Statistics* **67**(4), 435–466.

OECD (2004), Employment Protection Regulation and Labour Market Performance, *in* 'OECD Employment Outlook', OECD Publishing, Paris, pp. 61–125.

OECD (2010), *Measuring Globalisation: OECD Economic Globalisation Indicators*, Vol. 2009.
URL: *http://www.oecd-ilibrary.org/industry-and-services/measuring-globalisation-oecd-economic-globalisation-indicators-2005_9789264012394-en*

OECD (2011), *Divided We Stand: Why Inequality Keeps Rising*, OECD Publishing, Paris.

OECD (2015), *In It Together: Why Less Inequality Benefits All*, OECD Publishing, Paris.

OECD (2017), *OECD Employment Outlook 2017*, OECD Publishing, Paris.

OECD (2019), *Under Pressure: The Squeezed Middle Class*, OECD Publishing, Paris.

Oesch, D. (2006), 'Coming to grips with a changing class structure: An analysis of employment stratification in Britain, Germany, Sweden and Switzerland', *International Sociology* **21**(2), 263–288.

Oesch, D. (2013), *Occupational Change in Europe: How Technology and Education Transform the Job Structure*, Oxford University Press, Oxford.

Oesch, D. (2015), 'Welfare regimes and change in the employment structure: Britain, Denmark and Germany since 1990', *Journal of European Social Policy* **25**(1), 94–110.

Oesch, D. & Piccitto, G. (2019), 'The Polarization Myth: Occupational Upgrading in Germany, Spain, Sweden, and the UK, 1992-2015', *Work and Occupations* **46**(4), 441–469.

Oesch, D. & Rodriguez Menes, J. (2011), 'Upgrading or polarization? Occupational change in Britain, Germany, Spain and Switzerland, 1990–2008', *Socio-Economic Review* **9**, 503–531.

Parolin, Z. (2021), 'Automation, Occupational Earnings Trends, and the Moderating Role of Organized Labor', *Social Forces* **99**(3), 921–946.

Pissarides, C. A. (2001), 'Employment protection', *Labour Economics* **8**, 131–159.

Pontusson, J., Rueda, D. & Way, C. R. (2002), 'Comparative Political Economy of Wage Distribution: The Role of Partisanship and Labour Market Institutions', *British Journal of Political Science* **32**, 281–308.

Rendall, M. & Weiss, F. J. (2016), 'Employment polarization and the role of the apprenticeship system', *European Economic Review* **82**, 166–186.
URL: *http://dx.doi.org/10.1016/j.euroecorev.2015.11.004*

Ritchie, H., Roser, M., Mispy, J. & Ortiz-Ospina, E. (2018), 'Measuring progress towards the Sustainable Development Goals'.
URL: *https://sdg-tracker.org/*

Romer, P. M. (1994), 'The Origins of Endogenous Growth', *The Journal of Economic Perspectives* **8**(1), 3–22.

Ross, M. B. (2017), 'Routine-biased technical change: Panel evidence of task orientation and wage effects', *Labour Economics* **48**, 198–214.
URL: *http://dx.doi.org/10.1016/j.labeco.2017.08.007*

Rueda, D. & Pontusson, J. (2000), 'Wage Inequality and Varieties of Capitalism', *World Politics* **52**(3), 350–383.

Salvatori, A. (2018), 'The anatomy of job polarisation in the UK', *Journal for Labour Market Research* **52**(8), 1–15.
URL: *https://doi.org/10.1186/s12651-018-0242-z*

Salverda, W. & Checchi, D. (2015), Labor market institutions and the dispersion of wage earnings, *in* A. B. Atkinson & F. Bourguignon, eds, 'Handbook of Income Distribution', 1 edn, Vol. 2B, Elsevier B.V., pp. 1535–1727.
URL: *http://dx.doi.org/10.1016/B978-0-444-59429-7.00019-4*

Sebastian, R. (2018), 'Explaining job polarisation in Spain from a task perspective', *SERIEs* **9**(2), 215–248.
URL: *https://doi.org/10.1007/s13209-018-0177-1*

Sebastian, R. & Biagi, F. (2018), The Routine Biased Technical Change hypothesis: a critical review, Technical report, European Commission, Luxembourg.

Siegel, C. & Bárány, Z. L. (2018), 'Job Polarization and Structural Change', *American Economic Journal: Macroeconomics* **10**(1), 57 – 89.

Solow, R. M. (1957), 'Technical Change and the Aggregate Production Function', *The Review of Economics and Statstics* **39**(3), 312–320.

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.
URL: *http://www.journals.uchicago.edu/doi/10.1086/499972*

Tahlin, M. (2019), Polariseringsmyten - Försvinner verkligen de medelkvalificerade jobben?, Technical report, Arena Idé, Stockholm.

Thewissen, S. & Rueda, D. (2017), 'Automation and the Welfare State: Technological Change as a Determinant of Redistribution Preferences', *Comparative Political Studies* **52**(2), 171–208.

Thoening, M. & Verdier, T. (2002), A Theory of Defensive Skill-based Innovation and Globalization.

Tinbergen, J. (1974), 'Substitution of Graduate by Other Labor', *Kyklos* **27**(2), 217–226.

Treiman, D. J. (1977), *Occupational prestige in comparative perspective*, Quantitative studies in social relations, Academic Press, New York ; London.

United Nations (2008), *International Standard Industrial Classification of All Economic Activities (ISIC), Rev.4*, Department of Economic and Social Affairs, New York.

van der Wiel, K. (2010), 'Better protected, better paid: Evidence on how employment protection affects wages', *Labour Economics* **17**, 16–26.
URL: *http://dx.doi.org/10.1016/j.labeco.2009.09.008*

VanHeuvelen, T. (2018), 'Moral economies or hidden talents? A longitudinal analysis of union decline and wage inequality, 1973-2015', *Social Forces* **97**(2), 495–530.

VanHeuvelen, T. (2020), 'The right to work, power resources, and economic inequality', *American Journal of Sociology* **125**(5), 1255–1302.

Vincent-Lancrin, S. (2008), The Reversal of Gender Inequalities in Higher Education, *in* 'Higher Education to 2030: Volume 1: Demography', OECD Publishing, pp. 265–298.

Violante, G. L. (2008), Skill-Biased Technical Change, *in* S. N. Durlauf & L. E. Blume, eds, 'The New Palgrave Dictionary of Economics', 2nd edn, Palgrave Macmillan.

Visser, J. (2015), 'Data Base on Institutional Characteristics of Trade Unions, Wage Setting, State Intervention and Social Pacts, 1960-2014 (ICTWSS)'.

Visser, J. & Checchi, D. (2011), Inequality and the Labor Market: Unions, *in* B. Nolan, W. Salverda & T. M. Smeeding, eds, 'The Oxford Handbook of Economic Inequality', Oxford University Press, Oxford, chapter 10, pp. 1–48.

Vivarelli, M. (2014), 'Innovation, Employment and Skills in Advanced and Developing Countries: A Survey of Economic Literature', *Journal of Economic Issues* **48**(1), 123–154.
URL: *http://www.tandfonline.com/doi/full/10.2753/JEI0021-3624480106*

Wajcman, J. (2017), 'Automation: is it really different this time?', *British Journal of Sociology* **68**(1), 119–127.

Weisstanner, D. (2020), 'Insiders under pressure: Flexibilization at the margins and wage inequality', *Journal of Social Policy* pp. 1–20.

Western, B. & Rosenfeld, J. (2011), 'Unions, Norms, and the Rise in U.S. Wage Inequality', *American Sociological Review* **76**(4), 513–537.

White, R. C. (1931), 'Technological Unemployment', *Social Forces* **9**(4), 572–581.

Williams, M. (2017), 'Occupational Stratification in Contemporary Britain: Occupational Class and the Wage Structure in the Wake of the Great Recession', *Sociology* **51**(6), 1299–1317.

Wilmers, N. (2017), 'Labor unions as activist organizations: A union power approach to estimating union wage effects', *Social Forces* **95**(4), 1451–1478.

Winkler, O. W. (2009), *Interpreting Economic and Social Data - A Foundation of Descriptive Statistics*, Springer, Berlin.

World Bank (2021), 'World Development Indicators'.
URL:
https://databank.worldbank.org/reports.aspx?source=world-development-indicators#

Wright, E. O. & Dwyer, R. E. (2003), 'The patterns of job expansions in the USA: a comparison of the 1960s and 1990s', *Socio-Economic Review* **1**, 289–325.