WARWICK
THE UNIVERSITY OF WARWICK

A Thesis Submitted for the Degree of PhD at the University of Warwick

Permanent WRAP URL:
http://wrap.warwick.ac.uk/112180

Copyright and reuse:
This thesis is made available online and is protected by original copyright.
Please scroll down to view the document itself.
Please refer to the repository record for this item for information to help you to cite it.
Our policy information is available from the repository home page.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

warwick.ac.uk/lib-publications

Three Essays on the Impact of Economic Change on the Labour Market

by

Timo Leidecker

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Economics

February 2018

Contents

List of Tables

List of Figures

Acknowledgments

I would like to thank my supervisors Vera Troeger and Jennifer Smith for their continuous support and guidance for the entire duration of my PhD. I would also like to thank the faculty members at the Department of Economics, University of Warwick, for their helpful comments and suggestions whenever I needed them. Moreover, I would like to thank my friends at Warwick, who helped me in the process of writing this thesis with their much appreciated support, insights and company: Anik Ashraf, Chara Tzanetaki, Diego Solorazano, Horng Cherng, and Adam Hutchinson; Sakari Lemola, Irina Schumski, and Saleem Zuberi. My parents Elke Heidemann-Leidecker and Guido Leidecker deserve special thanks for providing me with the opportunity and support to pursue my studies for many years. Most of all I would like to express my love and gratitude for my partner Ayten Bilgin. It is difficult to imagine how I could have written this thesis without her support; and being with her is the most important result of this PhD.

Finally, I would also like to thank the ESRC for the financial support I received during the PhD.

Declarations

All the chapters in this thesis contain original research work. They have not been submitted by me for any other assessments or previous degree courses. All three chapters are solely my own work, which incorporates comments and suggestions from my thesis supervisors, Vera E. Troeger and Jennifer C. Smith, as well as from others I have discussed this work with. Chapters 1 and 2 contain figures based on LFS data provided to me by Jennifer C. Smith. For the derivation of variables in chapters 1 and 2, I also incorporated Stata codes, for consistent proportional mapping, provided by Jennifer C. Smith.

Abstract

This thesis consists of three chapters addressing three different yet related issues on the impact of economic change on labour markets.

In chapter 2, I assess the impact of United Kingdom (UK) job polarization at the worker-level by examining changes in the underlying labour reallocation. I use an annual random sample of UK employees from 1975 to 2015, based on NESPD and ASHE, following workers employed in a PAYE-registered job. To abstract from compositional changes, I conduct the analysis at the group level, distinguishing three age and gender groups. First, I identify distributional changes accounting for aggregate job polarization by decomposing employment share changes for low, medium, and high skilled employment into distributional and compositional changes. Second, I conduct a counterfactual exercise for changes in transition rates to compute their contribution to job polarization at the group and aggregate level. I find job polarization to be associated with a negative impact on young workers, who become more likely to start their career in low skilled jobs, and male workers, who experience longer non-employment periods. These changes combined can account for at least two thirds of the decline in the aggregate medium skilled employment share. Reallocation between job types appears unimportant.

In chapter 3, I examine changes in the distribution of non-employment spell durations associated with job polarization. I estimate the duration distribution in terms of survival functions, considering all exits to employment. I suggest a competing risks model allowing to decompose changes in survival functions into changes in hazard rates to low, medium, and high skilled jobs. Based on findings from chapter 2, I argue that changes in the hazard rate to medium skilled jobs are associated with job polarization. Survival functions are estimated non-parametrically for flow samples, based on NESPD and ASHE, of UK workers of six demographic groups entering non-employment in successive expansionary periods from 1975 to 2015. To organize the discussion, I distinguish short-term, longer temporary and permanent spells, finding that job polarization is associated with a general shift towards longer temporary spells, suggestive of longer reallocation periods, and male workers also becoming more likely to be permanently jobless, suggestive of a failure to reallocate. Women experience no comparable distributional changes, suggesting results are driven by aggregate and group-specific factors.

In chapter 4, I test whether skill-biased technological change (SBTC) differs across OECD countries. SBTC is often held to be an exogenous shock common to developed countries. I argue that seminal contributions establishing SBTC do not assess comparative aspects. Extending the approach by Katz and Murphy [1992] to a cross-country context, I test for SBTC differences using annual country-level panel data for 14 OECD countries from 1986 to 2010. I find evidence for significant variation. I explore whether differences are systematically related to institutional measures, for which I find tentative evidence.

Chapter 1

Introduction

For the alleged commodity 'labor power' cannot be shoved about, used indiscriminately, or even left unused, without affecting also the human individual who happens to be the bearer of this peculiar commodity. In disposing of a man's labor power the system would, incidentally, dispose of the physical, psychological, and moral entity 'man' attached to the tag.

The Great Transformation:
The Political and Economic Origins of Our Time
Karl Polanyi, 1944

The introducing quote by Karl Polanyi alludes to a distinguishing feature of labour economics: the subject of study, the labour market, differs from other markets as the commodities exchanged are, by necessity, inseparable from the men and women supplying their labour. This feature lends importance to the study of the impact of economic change on labour markets. It is not only that labour markets warrant a particular perspective because they function differently from other markets, or, say, because we are inherently more concerned about unemployed workers than excess supply of capital. It is the fact that the economic process is essentially dynamic, necessitating an ongoing reallocation of factor inputs, including labour, which emphasizes an additional aspect of labour markets: more than other factor inputs, labour may be slow, or even reluctant, to adjust to changes in the economic environment.[1]

[1] The notion that the economic process is essentially dynamic has been famously argued for by Schumpeter: 'Capitalism (...) is by nature a form or method of economic change and not only never is but never can be stationary (...). The opening up of new markets, foreign or domestic, and the organizational development from the craft shop and factory to such concerns as U.S. Steel illustrate the same process of industrial mutation (...) that incessantly revolutionizes the economic structure from *within*, incessantly destroying the old one, incessantly creating a new one. This process of Creative Destruction is the essential fact about capitalism.' (Schumpeter [2011], p. 82-83, emphasis

It stands to reason that workers are slow to adjust because they are not just suppliers of labour. As human beings, they respond to incentives other than price signals. A worker may be reluctant to leave a region with low employment prospects if this requires her to leave her elderly parents behind. A worker whose profession has become redundant may be slow to adapt, and consider retraining, because his identity and self-esteem are bound up with the job he used to hold. In fact, it is concern about this potential conflict, between the demands imposed by labour markets and the demands imposed by virtue of being human, and the ramifications this conflict may breed, which motivated Polanyi to write the book from which I quoted above.

Economic change, for instance in the form of trade intensification or the advent of new technologies, can therefore have a big impact on workers. It affects wages, destroys redundant jobs and creates new types of employment, often requiring workers to adapt and reallocate to new jobs. Even though workers may loose their jobs in the process, they are of course often better off as a result. Over the long-run, increases in labour productivity raised wages and have not been accompanied by a general decline in employment. The notion that technological change renders workers permanently non-employed has therefore been dismissed as 'lump of labour' fallacy. However, the process of labour reallocation is not instantaneous and, from the worker's perspective, can be disruptive. The jobs created are likely to be different, for instance in terms of skill requirements, from the jobs destroyed. The general rise in wages and employment can hide the suffering experienced by individual workers who struggle to adapt. The functioning of the economy, and the welfare of workers, both depend crucially on the impact of economic change on workers, and the relatively smooth adaption of labour to economic change.

These more general considerations motivate, in light of recent changes in labour markets, the focus of this PhD thesis. In Western countries, recent decades evidence a profound rise in wage inequality and changes in the employment distribution, both of which have been linked to technological change and trade intensification.[2] The earlier literature focused on rising skill premia, as wages of skilled workers, e.g. workers with at least some college education, outpaced those of unskilled workers. This rise in wage inequality occurred despite simultaneous increases in the skill endowment of the workforce. Explanations for this development focus on

in original).

[2]Seminal contributions to this literature are Bound and Johnson [1992]; Katz and Murphy [1992]; Card and DiNardo [2002]; Autor et al. [2003]. See Acemoglu and Autor [2011] for a more recent review of the literature.

trade, technology, and institutional change. Countries abundant in skilled labour, opening up to trade with countries abundant in unskilled labour, would shift their production towards industries using skilled workers relatively more intensively. The resulting rise in demand for skilled workers would push their wages up. Alternatively, technology complementing skilled workers relatively more than unskilled workers, referred to as skill-biased technological change (SBTC), would raise the relative productivity of skilled to unskilled workers. If workers are imperfect substitutes, this leads to an increase in relative wages. In addition to these factors, some authors argue that institutional change, the decline in union coverage for instance, would weaken wage compression, and thereby contribute to rising inequality.

Building on these findings, the recent literature suggests a more nuanced explanation, focusing on changes in the employment distribution: many countries see employment in jobs intensive in routine tasks, sometimes referred to as medium skilled jobs, disappear. Employment in such jobs declines, in absolute and relative terms, as employment in jobs intensive in manual or cognitive non-routine tasks, referred to as low and high skilled jobs, rises. These changes in the employment distribution are referred to as job polarization. Explanations for job polarization point to routine-biased technological change (RBCT): advances in many digital technologies imply that tasks which can be codified relatively easily, i.e. manual tasks, can be performed abroad, or replaced by capital, more and more cheaply. These shifts affect workers of various skill types differently, as it is argued that more skilled workers possess a comparative advantage in more complex, i.e. cognitive non-routine, tasks. Offshoring or capital accumulation substitute for routine tasks, acting to decrease labour productivity in medium skilled jobs. Workers shift to the extremes of the employment distributions as a result, with more skilled workers becoming more likely to hold high skilled jobs, and less skilled workers becoming more likely to hold low skilled jobs.[3]

The impact on the labour market has thus been profound, and the fact that workers of different skill types are differently affected implies that, despite rising wages and employment, many workers may be worse off as a result. Brynjolfsson and McAfee aptly summarize the problem: '[T]here's never been a better time to be a worker with special skills or the right education, because these people can use technology to create and capture value. However, there's never been a worse time to be a worker with only 'ordinary' skills and abilities to offer, because computers, robots, and other digital technologies are acquiring these skills and abilities at an

[3]See Autor et al. [2003]; Goos and Manning [2007]; Goos et al. [2009, 2014]; Acemoglu and Autor [2011].

extraordinary rate' (Brynjolfsson and McAfee [2014], page 11). Job polarization implies that workers have indeed become more likely to either hold a 'lovely' or a 'lousy' job.[4]

What then do we know about the impact of these changes in the economic environment on labour markets? The more recent literature has mainly focused on examining changes in the employment distribution. These distributional changes, reflecting job polarization, are themselves suggestive that workers have been impacted: workers are less likely to hold medium skilled jobs, and more likely to hold either low or high skilled jobs. However, while aggregate distributional changes show the fate of workers has changed, they do not show for whom it changed, and how. These questions need to be addressed in terms of labour reallocation: what happens to workers, who previously worked in medium skilled jobs, as these jobs disappear? Equivalently, how does labour, freed up from redundant jobs, reallocate to more productive uses? Distributional changes do not address these questions because they are compatible with several worker transitions. For instance, does the disappearance of medium skilled jobs reflect workers directly reallocating from medium skilled jobs to low or high skilled ones? Does the reallocation away from medium skilled jobs involve periods of non-employment? Do some workers fail to reallocate? Does it reflect older workers previously holding medium skilled jobs moving to new job types, or new cohorts entering the labour market moving directly to low or high skilled jobs?

The literature focusing on distributional changes therefore only provides limited insight into the impact of job polarization at the worker level. A small number of studies address these questions as they examine the process of labour reallocation underlying job polarization. However, these studies focus exclusively on US job polarization, and, due to data limitations, only consider a subset of possible worker transitions. US job polarization, these findings suggest, is indeed disruptive, involving transitions to non-employment or longer non-employment spells: medium skilled jobs disappear as workers previously holding such jobs move to non-employment.[5] While these findings add to our knowledge about the impact of job polarization in the US, the process of labour reallocation underlying job polarization in the UK has hitherto not been examined.

Chapter 1 attempts to fill this gap and provide evidence for changes in the process of labour reallocation accounting for job polarization in the UK. Doing so, I

[4] These terms have been coined by Goos and Manning [2007].

[5] See Smith [2013]; Cortes et al. [2014, 2016]; Foote and Ryan [2014].

provide some evidence on the impact of UK job polarization at the worker level. In particular, I use UK worker level data based on combining the New Earnings Survey Panel Dataset (NESPD) and Annual Survey of Hours and Earnings (ASHE), covering the period from 1975 to 2015, to examine which changes in transition rates importantly contribute to polarization. Contributions of transition rate changes are assessed using a counterfactual exercise similar to Cortes et al. [2014] and Shimer [2012]: holding particular transition rates constant at their level in an early baseline period, this exercise allows computing counterfactual changes in employment shares for low, medium, and high skilled jobs which would have pertained had this particular rate not changed, while all other rates adjusted. Because distributional or transition rate changes at the aggregate level can reflect changes in the composition of the workforce, which are not directly related to changes in outcomes at the worker level, I conduct the analysis at the group level, distinguishing three age and two gender groups. To identify how distributional changes for these groups contribute to aggregate polarization, I first conduct a decomposition of employment share changes into distributional and compositional changes. The counterfactual analysis then examines labour reallocation underlying polarization at the group level, while allowing to assess how these changes contribute to aggregate polarization.

I find that UK job polarization is largely associated with shifts towards non-employment: over the sample period from 1975 to 2015, male workers are less likely to hold medium skilled jobs at any point in time as they experience longer non-employment spells. This shift towards non-employment is driven by a decline in the non-employment outflow rate to medium skilled jobs, compatible with a decline in job creation for medium skilled jobs. Additionally, young workers entering the labour market for the first time are more likely to directly move to a low skilled job, rather than medium skilled jobs. These changes can account for the bulk of job polarization. Reallocation between job types, on the other hand, seems unimportant. These results have to be seen in light of possible mismeasurement of non-employment, but are largely compatible with additional evidence unaffected by mismeasurement. They are also in line with findings from the US. The narrative whereby Creative Destruction replaces redundant jobs with new ones and workers reallocate quickly, is not compatible with these results. The reallocation of labour underlying job polarization, these findings suggest, is indeed disruptive, and young and male workers have been affected most strongly.

Chapter 2 finds that reallocation related to job polarization involves longer non-employment spells: workers take longer time to move to new employment, possibly because fewer medium skilled jobs are being created. How long does it take workers to find new jobs? And do some workers fail to reallocate? The decline in non-employment outflow rates to medium skilled jobs suggest an increase in average non-employment duration, but these changes do not reveal whether all workers take longer time to find new jobs, or whether some workers leave permanently. Because US studies follow workers only for relatively short periods, they cannot assess how job polarization affects non-employment duration. These questions have therefore not yet been addressed in the context of job polarization.

This gap is addressed in chapter 3. NESPD and ASHE follow workers each year they are employed in a Pay-As-You-Earn (PAYE) registered job. As the sample runs from 1975 to 2015, workers can be followed for up to 40 years. While non-employment is not directly observed, with some restrictions the sample structure allows inferring non-employment from sample disappearance for workers aged 18 to 65 years, providing information also on long non-employment spells. I use this information to examine changes in the duration distribution associated with job polarization as follows: the duration distribution can be derived from a survival function considering exits to employment. For each duration, the survival function gives the probability, or fraction, of non-employed workers experiencing a non-employment spell lasting at least as long. Examining changes in survival functions for workers entering non-employment during successive expansionary periods, I ask which of these changes can be associated with job polarization. To address this question, I suggest a competing risks model which allows decomposing changes in the survival function into changes in hazard rates to low, medium, or high skilled jobs. Survival and hazard functions can be estimated non-parametrically using interval-censored observations. Again, I conduct the analysis at the group level to identify whether some groups have been affected disproportionately. As chapter 2 finds that polarization is associated with a decline in non-employment outflow rates to medium skilled jobs, I argue that changes in the duration distribution reflecting changes in hazard rates to medium skilled jobs can be associated with job polarization. To organize the discussion, I distinguish short-term (< 1 year), longer temporary ($1 \leq$ duration < 5 years) and permanent non-employment spells (≥ 5 years), and argue to interpret shifts to longer temporary spells as suggesting longer reallocation periods, while shifts to permanent spells suggest a failure to reallocate.

I find that job polarization is associated with a general shift to longer temporary spells, affecting all demographic groups. Beyond the decline in short-term spells, however, results differ substantially across groups. Male workers seem to have been more adversely affected. They generally experience larger shifts towards longer spells, and, importantly, a substantial increase in the probability to become permanently jobless. Women, on the other hand, have been affected less strongly, and experience some tendencies towards shorter spells in more recent years. While possible mismeasurement of non-employment again suggests interpreting these findings with caution, results are compatible with job polarization having a disproportionate negative impact on some, predominantly male, workers by rendering them permanently jobless. Job reallocation appears not only to be disruptive, as reallocation appears to generally take longer time, but, on this account, some workers appear to fail reallocating altogether. Differences across demographic groups suggest that these changes in the duration distribution, and the rise in permanent joblessness in particular, may reflect both aggregate and group-specific, arguably behavioural, factors.

Chapters 2 and 3 provide evidence for the process of labour reallocation underlying job polarization in the UK. As such, they provide some insight into how disruptive this process has been for workers. Findings suggest that labour reallocation is indeed disruptive, involving generally longer non-employment spells rather than direct reallocation between job types, and for some workers even results in permanent joblessness. The overall growth in employment appears to hide the suffering of some individual workers struggling to adapt. These findings therefore identify, and emphasize, issues to be addressed by policy makers: to facilitate labour reallocation and avoid permanent joblessness, or provide appropriate support. Differences across groups also inform future research questions which need to be addressed to guide policy making: for instance, as I do not observe the educational status of workers, I cannot condition results on educational groups. It stands to reason that reallocation patterns differ greatly among skill groups, which would have important implications for policy makers. If shifts towards non-employment occur mostly among relatively unskilled workers, one could improve the reallocation process by providing these workers with additional qualifications. Alternatively, reallocation may be slowed down because workers are reluctant to move to places with better employment prospects. If so, one could improve the process of reallocation by providing incentives for job creation in the affected areas, or support workers moving to find employment. Results from chapter 3 are especially suggestive: if longer non-employment durations reflect a decline in medium skilled job creation, why are women not more likely to become permanently jobless? If the rise in permanent joblessness reflects behavioural factors, answering this question would be informative about possible interventions aimed at avoiding permanent joblessness for male workers.

Concern for the supposed need of policy makers to affect the impact of technological change on the labour market also motivates chapter 4. The earlier literature on SBTC is often marked by what may be considered a fatalistic conception for the scope of policy makers to affect the impact of technological change on the labour market. The Krugman hypothesis famously argues that, in light of SBTC, policy makers have to choose between either rising wage inequality or higher unemployment (Krugman [1994]). Relatedly, the differential rise in unemployment during the 80s and 90s in the US and many European countries has been explained by interactions between common SBTC shocks and labour market institutions. SBTC would raise the productivity differential of skilled and unskilled workers, necessitating adjustment in either prices or quantities. In countries with less rigid labour market institutions adjustment took place in terms of prices, raising wage inequality. In countries with more rigid institutions, SBTC instead resulted in higher unemployment, as institutions introduced a lower bound for unskilled wages. The view that policy makers face a trade-off between inequality and unemployment is directly related to an underlying assumption about the nature of SBTC: SBTC is taken to be an exogenous shock common to developed countries. SBTC is thus assumed to be independent of country-specific factors that may be under the influence of policy makers.

In chapter 4, I argue that the underlying assumption of common SBTC is in need of empirical validation, and that the literature establishing SBTC as an international phenomenon fails to establish that SBTC exhibits the same direction and magnitude across countries. In fact, findings in the literature suggest variation in SBTC across Organization of Economic Co-operation and Development (OECD) countries, but often suffer from other shortcomings such as failing to control for institutional change affecting wage inequality. Using annual country-level panel data covering 14 OECD countries for the period from 1986 to 2010, I extend the major approach by Katz and Murphy [1992] to a cross-country context. I address the identified shortcomings by including additional controls. In particular, I use factor variables derived from various institutional measures as proxies for the institutional environment. I first use this approach to test whether SBTC differs across countries. Finding significant differences, I examine possible factors accounting for the cross-country variation. Important factors varying at the country-level are institutions, and some authors have indeed argued that institutional rigidity can induce firms to raise the productivity of unskilled workers. Second, I therefore ask whether SBTC differences are systematically related to country-specific measures of the institutional environment. I find tentative evidence that SBTC varies with institutions, with some institutional measures being related to a smaller skill-bias of technological change.

Although the level of analysis is not well suited to identify particular mechanisms, results suggest that policy makers can directly affect technological change, and so its impact on the labour market. This finding is suggestive for future research to examine possible mechanisms by which policy makers can do so.

The work as a whole contributes to our understanding of the impact of economic change on the labour market in terms of the impact of job polarization at the worker-level, and of the scope of policy makers for affecting the impact of technological change in terms of SBTC. More detailed conclusions are presented at the end of each chapter.

Chapter 2

How Disruptive Was Labour Polarization for UK Workers?

2.1 Introduction

Recent events, from electoral results in Western countries to talk of the jobless economy, evidence an increasing concern about the effects of economic change, i.e. technological change and trade, on employment prospects of particular worker groups. Concern about the impact of technological change on the labour market is not new. Going back at least as far as the Luddites, fear has been that improvements in labour productivity harm workers by rendering their source of income, their labour, redundant. However, as ever increasing labour productivity has not been accompanied by a decrease in aggregate labour input, this fear has been repudiated as being based on faulty theory, the "lump of labour fallacy":[1] rather than decreasing the total number of jobs, increasing labour productivity reflects efficiency enhancing labour reallocation, i.e labour in unproductive jobs is freed up to be used for more productive purposes.

Yet, recent changes in the labour market suggest it is important to look beyond the total number of jobs. In recent decades the impact of technological change and trade intensification on the labour market has been summarized by the concept of job polarization. As jobs intensive in routine tasks can be more easily off-shored or replaced by machines, the employment distribution in many OECD countries has shifted away from medium skilled towards low and high skilled jobs.[2] Changes in the employment distribution are of relevance from the worker's

[1] See Autor [2014] for a discussion of automation anxiety and the lump of labour fallacy.

[2] Goos and Manning [2007]; Salvatori [2015]; Holmes and Mayhew [2012] document job polarization for the UK. Autor et al. [189-194] document job polarization for the US. Goos et al. [2009]

perspective because jobs are different. They can pay high or low wages, the work can be rewarding or dull, physically challenging or mentally stressful. Our social rank and – as many academics may testify – the very source of our identity can rest on the job we hold. In short, jobs are different and can be 'lovely or lousy', and as workers move from one type of job to another they can be better or worse off. Job polarization implies workers are now more likely to hold either a 'lovely' or a 'lousy' job.[3]

While distributional changes are suggestive that workers can be better or worse off, it does not tell us how disruptive economic changes associated with these distributional shifts have been. Job polarization, changing the types of jobs workers hold, may have been disruptive for some and beneficial for other workers. Overall, the impact of job polarization at the worker level remains unclear.

What determines these costs? Underlying these distributional changes is the process of Creative Destruction. Changes in technology and trade render some jobs obsolete but create opportunities for new ones. How workers are affected by this process, and so how disruptive creative destruction is, depends on what happens to workers whose redundant jobs are destroyed. How does labour, freed up from redundant jobs, reallocate to more productive uses? Distributional changes show the fate of workers has changed, but they do not show for whom it changed, and how. For instance, they do not reveal whether workers whose medium skilled jobs were destroyed went to low or high skilled jobs, or how long it took to be reallocated, or whether they left the labour market altogether. To understand how disruptive job polarization has been, one needs to focus on changes in labour reallocation, i.e. changes of worker transitions associated with job polarization.

This is the aim of this paper. It examines the question how disruptive job polarization has been in the UK from the worker's perspective. It does so using the following analyses.

First, I identify distributional changes at the level of demographic groups which account for aggregate job polarization. Distributional changes and changes in

and Acemoglu and Autor [2011] show that job polarization is present in many European countries. Trade and technological change have been established as the main determinants of job polarization. Autor et al. [2003] first suggested the routinization hypothesis, according to which recent technologies complement non-routine and substitute for routine tasks. Michaels et al. [2014] use industry-level data for 11 OECD countries to show that ICT is related to shifts in demand from medium to high skilled workers. Using a theoretically motivated shift-share analysis, Goos et al. [2014] find that offshoring and routine-biased technological change (RBTC) can explain most of job polarization experienced in a cross-country context, focusing on West European countries. More recently, Autor et al. [2015] use detailed commuting zone data for the US to identify the differential effect of exposure to trade or RBTC on the employment structure, finding that both act to decrease the share of routine-intensive occupations.

[3] See also Goos and Manning [2007].

labour reallocation at the aggregate level may reflect compositional changes, rather than changes in worker outcomes. The focus at the demographic group level serves to abstract from such compositional changes in the workforce, and it additionally allows investigating to what extent worker groups have been disproportionately affected. I do so by decomposing changes in aggregate employment shares into distributional and compositional changes at the demographic group level, considering six demographic groups (two gender and three age groups). Overall, this analysis serves to identify changes in worker outcomes associated with job polarization, which then need to be explained in terms of changes in labour reallocation.

In the second part of the paper, I ask how changes in labour reallocation relate to these changes in employment shares. Measuring job reallocation by annual worker transition rates, I use a counterfactual analysis to examine how changes in transition rates during expansionary periods are associated with job polarization. Again I abstract from compositional changes by conducting the analysis at the demographic group level. Results of the counterfactual analysis are suggestive for the importance of changes in particular transition rates during expansions, and so labour reallocation over the long-term, in accounting for job polarization. Finally, this allows me to discuss how workers have been impacted by job polarization.

Examining the relationship between job reallocation and job polarization, and thus understanding how disruptive job polarization has been, is important as it tells to what extent job polarization had adverse effects on workers, how workers were affected, and who has been most affected. Giving answers to these questions is relevant for policy makers as it provides information on the potential magnitude of any adverse impact on workers, how to target policies to support the workers most affected, and on proximate reasons for workers being worse off, e.g. whether job polarization was driven by job creation or job destruction.

Results suggest the potential for job polarization being very costly for some worker groups. There is little evidence that job polarization occurs as workers directly reallocate from medium skilled jobs to low or high skilled employment. There is no evidence that polarization is generally associated with workers upgrading to high skilled jobs, which would constitute workers benefiting from job polarization. Instead, polarization mostly reflects workers reallocating from medium skilled jobs to non-employment as they no longer return to medium skilled jobs. Although one has to be cautious about the measurement of non-employment, job polarization appears to largely reflect shifts from medium skilled jobs to non-employment, driven by a decrease in job creation rather than an increase in job destruction. While these shifts appear to be most important for prime aged and older male workers, young men and women directly enter into low rather than medium skilled jobs, thus shifting from medium to low skilled employment. Prime aged and older women seem unimportant to account for past job polarization, but are likely to gain in importance as they comprise a larger share of the workforce. Overall, this implies that job polarization may have been very disruptive for some workers, resulting in longer non-employment or workers starting their career in worse jobs.

This paper adds to the literature on job polarization and labour reallocation by describing the possible impact job polarization has on workers. Several papers have examined the link between worker transition rates and job polarization in the US. The most closely related paper is Cortes et al. [2014]. Cortes et al. [2014] conduct a counterfactual analysis using US data to examine the importance of aggregate transition rates for the decline in routine cognitive and manual employment shares. They mostly focus on transitions from and to unemployment and non-participation. Results indicate that the combined decline in transition rates from unemployment to routine jobs, and from non-participation to routine jobs, and to lesser extent higher rates from routine employment to non-participation, can account for close to two thirds of the decline in routine employment. Most of this appears to be concentrated in the period following the Great Recession, however. Cortes et al. [2016] conduct a decomposition of shares in labour market states in routine and non-routine cognitive and manual employment, as well as non-employment, using the same data. They find that most of the decline reflects distributional changes for demographic groups, and they identify young and prime-aged male workers with low levels of education and young and prime-aged female workers with medium levels of education accounting for most of the overall decline. They also find that these workers largely shifted to non-routine manual jobs and non-employment.[4]

An additional contribution comes from Foote and Ryan [2014] who focus on cyclical changes in transition rates related to US job polarization. Their analysis differs from mine, apart from the focus, in that it is based on a multinomial logit to estimate transition probabilities, rather than conducting a counterfactual analysis to examine the importance of transition rate changes for the decline in medium skilled jobs. Findings indicate that unemployed medium skilled workers are most likely to return to medium skilled jobs, or to remain in unemployment, and both rates

[4]Another related paper is Smith [2013]. The author uses US data to compute counterfactual steady state shares for low, medium, high skilled occupations and others to examine which transition rate changes are important for job polarization. Results generally confirm Cortes et al. [2014] but differ to the extent Smith [2013] uses a shorter time period, reflecting inconsistencies in the dataset for the pre and post 1994-period, and a classification of workers into skill-categories based on broad occupational categories rather than mapping, based on Autor et al. [2003], detailed occupational codes into the routine, non-routine dimension.

vary over the business cycle. The probabilities to move to low or high skilled jobs are acyclical and small. Importantly, they find that more highly educated workers are somewhat more likely to move to high skilled jobs. They also provide evidence that groups subsequently experiencing more intensive job polarization are also more likely to drop out of the labour force.

All these studies examine job polarization in the US, using Current Population Survey (CPS) data and focusing on the period 1994 to 2012 or 2014. Cortes et al. [2014, 2016] use data before 1994 at the expense of not being able to reliably examine job-to-job transitions. To the best of my knowledge, no studies exist examining worker transitions underlying job polarization in the UK. Upward and Wright [2007] come closest, examining worker transitions underlying skill-upgrading in the UK and US. Using British Household Panel Survey (BHPS) data for the UK covering the period 1991 to 2001, they examine to what extent skill-upgrading at the plant-level occurs by upgrading existing low skilled workers or hiring new high skilled workers. They use probit models to estimate the probability to either downgrade or upgrade within and between firms, or to become unemployed, in the presence of technology shocks. Shocks are proxied by changes in industry-level shares of high skilled workers. They distinguish four skill groups based on International Standard Classifications (ISC88), which in turn are based on occupations' required level of general education and training. Results indicate that low skilled workers become more likely to upgrade when there is a technology shock, but overall workers are more likely to downgrade or exit to unemployment.

Related to the decomposition conducted in the first part of this paper is Salvatori [2015]. Using mainly UK Labour Force Survey (LFS) data covering the years 1979 to 2012, and using a similar classification of workers into skill categories, the author decomposes employment shares of occupational median wage deciles into changes within and between demographic groups. He finds the largest increase in employment shares has been for the top two deciles, which grew by 16pp, and this increase is entirely accounted for by the increasing share for high skilled workers. The decline in middling occupations of 19pp entirely reflects changes for non-graduates, two thirds of which reflect distributional changes, entirely directed towards lower skilled occupations. The remaining third would be explained by the declining share of non-graduates.

The paper also adds to the literature relating job reallocation to aggregate productivity growth.[5] The upshot of this literature is that productivity growth requires the reallocation of resources towards their most productive use. Creative

[5] See for instance Foster et al. [2001, 2008]; Haltiwanger [2011].

Destruction implies that some firms become more productive than others. Aggregate productivity growth requires that labour and capital move from less to more productive firms. The link between labour reallocation and productivity growth varies across countries, however, suggesting that country-specific factors, such as market structures or institutions, are important areas for interventions by policy makers. Reallocation therefore accrues benefits, but also imposes costs. Part of these costs are the adjustment costs borne by workers. Bartelsman [2013] points to potential future innovations, further aggravating the need to provide new jobs for displaced workers. Understanding how job polarization has affected labour reallocation in previous decades addresses this trade off between productivity growth and costs imposed on workers by providing evidence on the costs borne by workers.

The paper proceeds as follows. Section 2.2 describes the dataset and the variables used for this analysis as well as the limitations imposed by the dataset. The decomposition of population share changes is discussed in section 2.3. In section 2.4 I discuss the counterfactual analysis. Section 2.5 concludes and discusses results.

2.2 Data

In this section I briefly describe the dataset used for the analysis and the construction of main variables. More detailed explanations are given in the data appendix.

I combine NESPD and ASHE to create a combined dataset spanning the period 1975 to 2015.[6] Both datasets sample employee jobs based on the two last digits of their National Insurance number, resulting in a one percent random sample of jobs paying tax and National Insurance contributions via the PAYE system. Workers sampled based on their National Insurance number occur in the sample in all subsequent years in which they are employed in a PAYE registered job. PAYE is the UK system by which employers can collect tax on behalf of the UK government prior to paying workers. Jobs are sampled in January and information is retracted from employers with regard to a reference date in April. ASHE constitutes a continuation of NESPD, but improves on NESPD with regard to depicting low wage employment more accurately.[7]

I restrict the sample to workers aged 18 to 65. For workers who hold several jobs I only keep observations on their main job. Note that the sampling structure

[6] See ONS [2017d,a].

[7] The most relevant difference between ASHE and NESPD for the purpose of this analysis is that ASHE improves on NESPD by following up on non-responses. Non-responses can occur either due to job changes which are not reflected in PAYE at the time of the sample selection, or because they occur between sample selection and the survey reference date.

implies that workers occur in the sample whenever they are employed in a PAYE registered job. Workers who previously occurred in the sample may be missing in subsequent years for one of the following reasons: Non-employment, self-employment, emigration, or death. I abstract from the last three and refer to workers disappearing from the sample as non-employed in the UK.[8] The resulting dataset is an annual panel at the worker-level, observing the worker's wage, 3-digit occupational code, gender, age, and whether or not he or she has been employed at the reference date. Defining the working population as workers aged 18 to 65 years old, workers enter the sample the first time they are employed in an income tax paying job from 1975 onwards, and they remain in the sample until they reach retirement age.

The advantage of using NESPD and ASHE is data quality, as information is provided by employers, as well as the large sample size and long time period. While annual observations imply that high-frequency changes are not observed, year-to-year transitions lend themselves to the analysis of long-term employment trends.[9] Adding to this is the fact that workers remain in the sample as long as they are employed in a PAYE registered job. While the current analysis only makes use of year-to-year worker transitions, in chapter 3 I extend on this paper's findings, using survival analysis to exploit long-term information on worker transitions lasting up to 40 years.

The dataset poses the following challenges. The combined dataset exhibits discontinuities in 1990, 2001 and 2010 due to changes in occupational codes, in 1996 due to the change from NESPD to ASHE, and in 2006 due to the introduction of automatic occupation coding. These discontinuities cause jumps in transition rates. I disregard observations for years in which these discontinuities occur when constructing transition rates.[10] Additionally, the sampling procedure implies that non-employment stocks and transition rates out of non-employment in early years suffer from sample start problems. Most importantly, these may affect the change in transition rates from non-employment in the early years of the sample. I address these problems by additionally using transition rates corrected for the likely bias

[8] As my analysis compares variables over time, this abstraction affects results to the extent the prevalence of non-appearance due to any of these reasons changes over time.

[9] Annual observations imply non-employment inflow and outflow rates may suffer from time aggregation. Time aggregation implies that the level of transitions is understated. For instance, annual observations imply I do not observe workers who move in and out of non-employment in between observations, thereby understating the non-employment inflow and biasing the non-employment inflow rate downwards. The impact of time aggregation is typically discussed in the analysis of cyclical fluctuations in transition rates. See for instance Shimer [2012]. It is of minor importance for my analysis as I compare inflow rates between different expansionary periods.

[10] Consider, for instance, the change in occupational coding from 1990 to 1991. As new codes are used in 1991, transition rates, measuring the probability of workers changing their occupation from 1990 to 1991, exhibit jumps in 1990 as workers are more likely to change occupations.

introduced by the sample start problem. I also provide independent evidence compatible with results drawn from using these transition rates. I provide more details in the text whenever relevant.

2.2.1 Choice of Time Periods

As my goal is to examine long-term changes in worker's employment prospects, I focus on expansionary periods. Contractionary periods are themselves of interest for job polarization, but they constitute periods of short-lived, sharp shocks, which warrant particular attention. For this reason I leave the focus on recessions to future research. I delineate expansionary from contractionary periods based on quarterly Gross Domestic Product (GDP) data. Contractionary periods are identified in terms of declining GDP from peak to through.[11] Using quarterly data covering the sample period, contractionary periods from peak to through are: 1979q2-1981q1, 1990q2-1992q2, 2008q1-2009q2. Based on this inspection I choose as contractionary periods: 1979-1981, 1990-1992, and 2008-2009. Consequentially, for the decomposition expansionary periods are: 1975-1979, 1981-1990, 1992-2008, 2009-2015. This reflects the reasoning that the decomposition over an expansionary period should cover the entire period during which GDP was increasing.[12]

2.2.2 Construction of Skill Categories

I group workers into job types based on their 3-digit occupation. First, I use consistent proportional mapping, described in more detail in data appendix A.1, to obtain a variable containing workers' Standard Occupational Classification (SOC) 2010 codes covering the entire sample period. Second, similar to Goos and Manning [2007], I compute skill groups based on a time invariant ranking of 3-digit occupational median wages. Specifically, I compute median wages for each 3-digit occupation in 1975, based on gross weekly wages, and rank occupations into deciles according to their 1975 median wage.[13] Third, to map occupations to job categories

[11] I use ONS GDP ABMI quarterly time series data (ONS [2017c]). Alternatively, I also examined the annual time series ONS Employment Seasonally Adjusted (SA) 16-64 (ONS [2017b]), which yields the following periods as periods of declining employment: 1979-1983, 1990-1993, 2008-2010. Roughly the same periods are identified, but due to slower recovery of employment the periods tend to be longer. As expansionary periods are better reflected by GDP, and noting that either series suggest roughly the same choice of periods, I base the delineation on the GDP time series.

[12] For instance, for the expansionary period after the first recession, this period spans from the low-point of the recession in 1981 to the peak before the subsequent recession in 1990.

[13] Alternatively, one could use employment weighted deciles, as originally done by Goos and Manning [2007]. Results using employment weighted deciles are very similar for both the decomposition and counterfactual analysis. I use unweighted deciles because of the intuitive appeal of using occupations as the unit of analysis. Additionally, the implied ranking is independent of the initial

associated with job polarization, I examine employment share changes by median wage deciles over the sample period. Figure 2.1 shows percentage point (pp) changes of employment shares over the period 1975 to 2015. The figure clearly shows that job polarization is present over the sample period. Employment share changes are largely compatible with similar studies using LFS data. The main difference is that the rise in top deciles is less pronounced in figure 2.1.[14] I group occupations in the lowest (highest) deciles with growing employment share into the low (high) skilled employment category, and all occupations in middling deciles with declining employment shares are grouped into the category for medium skilled employment. Thus, in the subsequent analysis jobs in occupations in the 1st decile are referred to as low skilled employment, in the 9th and 10th deciles as high skilled employment and in the 2nd to 8th deciles as medium skilled employment.

In each year each worker is assigned to one of four mutually exclusive labour market states. If the worker occurs in the sample in the current year, based on the worker's occupation he or she is assigned to be employed in a low, medium, or high skilled job. If the worker was employed in earlier years but does not appear in the dataset in the current year, he or she is assigned to be non-employed, as long he or she is of working age.

Table A.1 in the appendix shows the composition of each skill category in terms of SOC 2010 major occupational groups for the full sample.[15] Low skilled employment comprises mostly service jobs, i.e. elementary services, sales and personal care, and office admin. These occupations comprise more than three fourths of low skilled jobs. Elementary services comprise tasks such as cleaning, protection, food production etcetera. High skilled jobs comprise mostly professionals and managers. These two occupations account for almost 90 percent of high skilled jobs. As expected, the majority of medium skilled jobs comprises jobs in office and admin; skilled trade; associates and technicians; and operators and production. These occupations comprise 68.8 percent of medium skilled jobs.[16]

---

employment distribution. This may be preferable because, as shown below, employment shares exhibit job polarization from the sample start, so the beginning of job polarization arguably predates the beginning of the sample.

[14] See, for example, Goos and Manning [2007]; Holmes and Mayhew [2012]; Salvatori [2015]. These conclusions remain unchanged when using employment weighted deciles. One explanation for the smaller rise in top deciles could be the omission of self-employed workers in NESPD and ASHE, assuming these workers have high earnings. Also, PAYE is said to exclude earners with more complex financial affairs, which tend to be high income earners.

[15] Occupations belonging to the same major occupational group can be assigned to different skill categories because classification is done at the 3-digit level.

[16] Perhaps surprising is the large share of elementary occupations for medium skilled jobs. One might think that elementary occupations would be grouped as low skilled employment if using a broader definition of low skilled employment, i.e. including more of the lower deciles. Interest-

Figure 2.1: Change in Employment Shares for Median Wage Deciles, 1975-2015

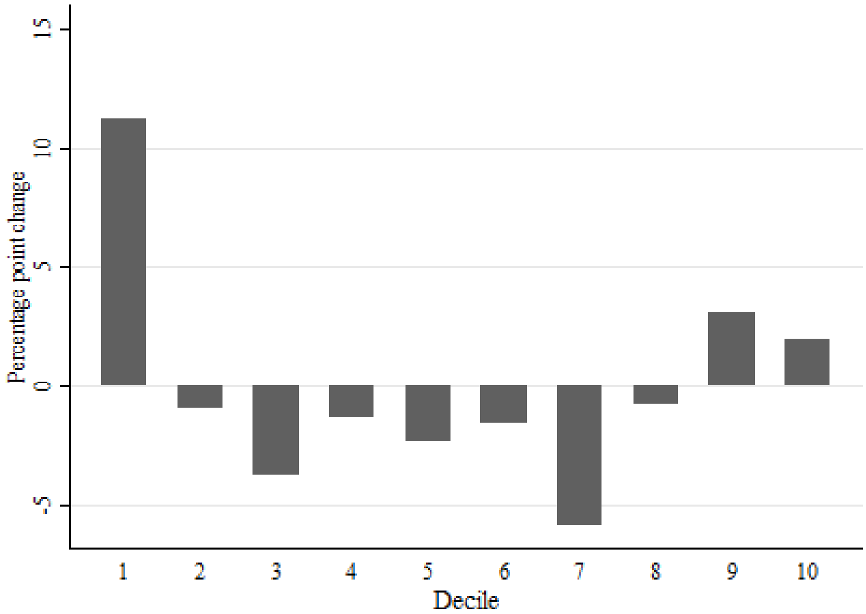

Employment share changes of median wage deciles over the period 1975 to 2015. Based on annual observations of NESPD and ASHE. Deciles are based on median wage of SOC 10 occupations at 3-digit level in 1975. Annual decile employment shares are constructed in 1975 and 2015 as the ratio of the number of employed workers in the respective decile group to total employment.

The resulting mapping of occupations into skill categories is also largely compatible with the suggested categorization of Acemoglu and Autor [2011], which relates to the distinction of low skilled jobs as intensive in routine and non-routine non-cognitive tasks, medium skilled jobs in routine cognitive tasks, and high skilled jobs in non-routine cognitive tasks. Using United States (US) SOC major occupational groups, the authors suggest a classification of major occupational groups into three skill categories based on employment growth over the period 1979 to 2009. Managerial, professional and technical occupations are classified as high skilled jobs. Sales; office and administrative support; production, craft and repair; operator, fabricator and labourer are classified as medium skilled jobs. Production, craft and repair and operator, fabricator and labourer are classified as medium skilled jobs despite positive employment growth over the respective period, as growth lagged behind average employment growth. Finally, they classify service occupations com-

ingly, however, it appears that most elementary occupations judged to be medium skilled are in middling deciles, i.e. the 3rd and higher deciles. Examples for elementary occupations classified as medium skilled are elementary construction occupations, elementary administrative occupations, and elementary storage occupations.

prising protective service, food preparation and cleaning services as well as personal care as low skilled jobs.

Figure A.1 in the appendix shows that the ranking of occupational median wages is largely time invariant, and so the choice of base year is inconsequential for the mapping of median wage deciles into skill categories. Importantly, occupations grouped into low (high) skilled employment earn the lowest (highest) median wage in 1975, 1995 and 2015. Figure A.2 in the appendix shows that changes in employment shares are generally associated with changes in absolute employment. Median wage deciles assigned to low and high skilled employment clearly show absolute employment growth over the sample period. Median wage deciles assigned to medium skilled employment largely show declining absolute employment. The 3rd and 8th deciles, which exhibit positive annual growth rates close to zero, exhibit near constant absolute but declining relative employment.

As demonstrated by figure 2.1, job polarization was clearly present from 1975 to 2015, the overall sample period. To see how polarization varied over time, figure 2.2 shows the annual time series of employment shares for each skill category. Periods identified below as recessionary periods from peak to through are indicated by dashed vertical lines. First, one can see that job polarization is present from the beginning of the sample period onwards, although at seemingly slower pace in earlier years. Second, most of the decline in medium skilled employment shares, definitive of job polarization, occurs during expansionary periods. This fact justifies examining the link between job polarization and labour reallocation focusing on expansionary periods, but this is not to say that recessions do not matter. Polarization continues during recessions, and even exhibits some comparatively sharp changes during downturns. Jaimovich and Siu [2012] examine the link between job polarization and recessions for the US. They argue that the decline in medium skilled jobs is concentrated in the aftermath of recessions, possibly as firms use recessions as opportunities to lay off redundant workers. As figure 2.2 shows, this is not generally true for the UK. Compatible with their findings, during the 80s the share of medium skilled jobs remains fairly constant throughout expansionary years, and falls more steeply around the early 1990s recession. On the other hand, medium skilled jobs continue to decline throughout the long, expansionary period starting during the 90s and ending with the Great Recession, and the decline continues unabated in all years following the Great Recession. Third, while the employment share of medium skilled employment decreases more or less constantly over the entire sample period, the growth of employment shares for low and high skilled employment changes pace several times over the sample period. Most strikingly, low and high skilled employ-

Figure 2.2: Annual Job type Employment Shares, 1975 to 2015

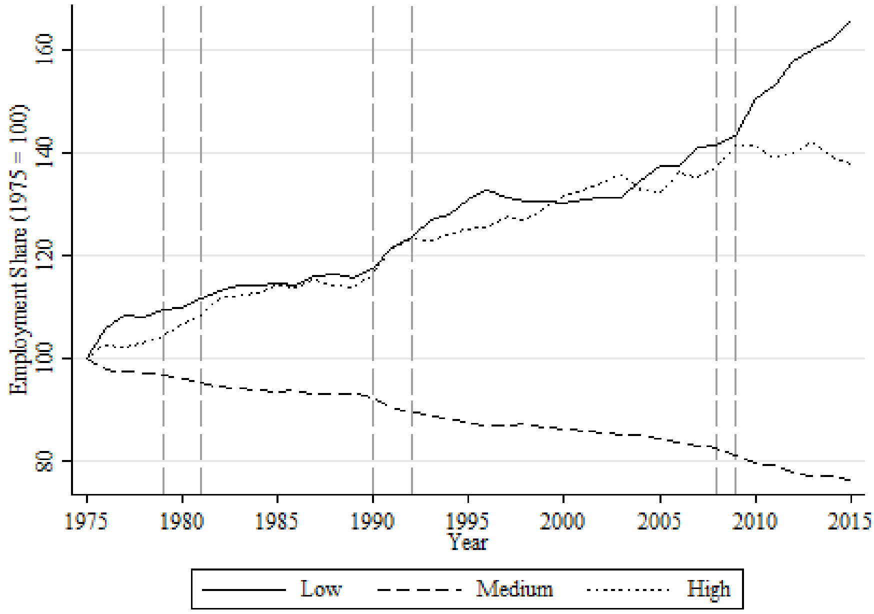

Standardized annual employment shares of low, medium and high skilled jobs from 1975 to 2015. Employment shares in 1975 are standardized to 100. Based on annual observations of NESPD and ASHE. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Vertical lines indicate recessions.

ment shares increase proportionally until the end of the Great Recession, while in the most recent period the share of high skilled jobs is levelling off, and the share of low skilled jobs is increasing at an unprecedented pace. This may suggest a growing importance in low skilled jobs for reallocation: the disappearance of medium skilled jobs is associated with workers shifting to low rather than high skilled jobs.

2.3 Job Polarization and Distributional Changes

The aim of this chapter is to examine the changes in job reallocation which underlie the shift of the employment distribution away from medium skilled jobs, known as job polarization. Job polarization is suggestive for changes in outcomes at the worker-level: employment in medium skilled jobs declines because some workers are less likely to hold a medium skilled job. Yet, not all of job polarization has to result from such changes. It can also reflect changes in the composition of the workforce. For instance, job polarization may result from an increasing participation rate among women, who are more likely than other worker groups to work in low or high skilled jobs. This section aims to abstract from such compositional changes and identify changes in worker outcomes which are associated with aggregate job polarization.

To distinguish changes in worker outcomes from changes in composition, I conduct a decomposition of employment share changes into compositional and propensity changes, distinguishing six demographic groups. I identify worker outcomes in terms of propensity changes for a given group, and I distinguish these from job polarization reflecting changing weights of demographic groups. The subsequent section then examines the identified worker outcomes in terms of labour reallocation.

I conduct the decomposition of employment shares for employment in low, medium, and high skilled jobs.[17] First, I restrict the sample to employed workers, and then group workers in the sample into men and women and three age groups (18-30, 31-50, and 51-65 years), giving six groups overall. Following Cortes et al. [2016], I employ the following decomposition:

$$\pi_1^j - \pi_0^j = \sum_{g=1}^{G} \Delta w_{g0} \pi_{g0}^j + \sum_{g=1}^{G} w_{g0} \Delta \pi_{g1}^j + \sum_{g=1}^{G} \Delta w_{g1} \Delta \pi_{g1}^j \tag{2.1}$$

where $\bar{\pi}_1^j$ is the aggregate employment share of employment type j at time 1, π_{g0}^j is the employment share of group g, out of a total of G groups, of employment type j at time 0, and w_{g1} is the employment share of group g at time 1.[18]

The first component in equation 2.1 refers to the composition effect, the second to the propensity effect, and the third is an interaction effect.[19] Composition, propensity and interaction effects refer to fractions of the total change of the respective employment share. In the tables below, I multiply fractions by 100 to get percentages.

Based on this decomposition, I address the following questions: first, how much of job polarization reflects changes in worker outcomes? This question is

[17] In an earlier version I conducted the decomposition and the counterfactual analysis for population shares of employment in low, medium, high skilled jobs and non-employment. Results suggest large shifts towards non-employment for all groups except young female workers. Counterfactual results are qualitatively unchanged. I choose to conduct the analysis for employment shares because the possible mismeasurement of non-employment complicates the interpretation of results.

[18] I conduct this decomposition over each expansionary and contractionary period, as well as for the period ranging from 1981 to 2015. I focus on results for this latter period to relate results to the counterfactual analysis, which will compare benchmark and counterfactual changes in employment shares over the same period, taking the first period, from 1975 to 1979, as reference for counterfactual transition rates.

[19] The interaction effect, at the group level, acts to increase the employment share if composition and propensity move in the same direction. For example, the interaction effect for prime aged male workers adds to the increase in the employment share of low skilled jobs if these workers are both becoming more likely to work in low skilled jobs, and growing in size relative to other groups.

addressed in section 2.3.1. Second, which changes in worker outcomes are associated with job polarization? Dividing this into 2 sub-questions, section 2.3.2 examines which group's propensity changes were important for job polarization, and section 2.3.3 discusses how groups' propensities changed. Overall, this serves to identify propensity changes which are examined in terms of labour reallocation in section 2.4.

2.3.1 Overall Decomposition

Table 2.1 shows results for the decomposition shown in equation 2.1. I refer to this as the overall decomposition as components are summed over all groups. The first and second columns give the labour market state and period for which the decomposition is conducted. The third column gives the respective employment share at the beginning of the period. Column four gives the total change. Columns five to seven give the percentage point change attributable to changes in the composition, propensity, and an interaction term.[20]

Focusing first on total period changes in employment shares in column four, one can see how table 2.1 mirrors the pattern shown in figure 2.2: job polarization is present throughout the sample period. Changes for each sub-period exhibit a similar pattern, with the share of medium skilled jobs declining and the share of low and high skilled jobs increasing. This holds true for all periods except the most recent one, when the share of low skilled employment grows much more strongly, and the share of high skilled jobs declines. Cumulative changes in employment shares are largest for expansionary periods, but polarization continues throughout recessions. Overall, this is compatible with job polarization occurring over the long-run, possibly exhibiting a change in pattern as shifts to low skilled jobs grow stronger, but showing no sign of a slowing down.

How can this pattern be explained in terms of propensity and compositional changes? The most important feature shown in 2.1 is that the decrease in the employment share of medium skilled jobs mostly reflects changes in propensity for given demographic groups. This generally holds for the overall period as well as sub-periods. The only exception is period 3, from 1981 to 1990, when the decline reflects compositional and propensity changes by equal amounts. Changes in low

[20] Note that the overall compositional or propensity effect is not equivalent to the sum of period-to-period effects. That is, if propensity changes occur from one period to the other, subsequent periods compute propensity and compositional effects based on the propensities prevailing at the end of the previous periods. Also, the interaction term reported in table 2.1 is summed over all groups. Because it multiplies group-specific composition and propensity changes, it is possible that the interaction term is negative even though the sums of composition and propensity terms are of the same sign.

and high skilled employment shares similarly result predominantly from changes in the probability of particular workers groups to hold either a low or high skilled job. Changes in the workforce contribute to the fall in medium skilled jobs and the rise in low skilled jobs, but counteract the rise in high skilled employment. This implies that workers entered the labour force who are more likely to work in a low skilled job, and less likely to work in either a medium or a high skilled job. The subsequent sub-section allows identifying which groups account for these shifts.

Salvatori [2015] also finds that job polarization in the UK occurred over the long-run. As he uses more detailed demographic groups, also controlling for education, his findings provide a useful benchmark for the scope of propensity changes. In particular, while he finds that the increase in employment shares of high skilled deciles is due to the increase in educational attainment, the decrease in medium skilled deciles is mostly due to distributional changes of non-graduates.

Overall, this shows that the relative and absolute disappearance of medium skilled jobs reflects changes in worker outcomes: workers for given demographic groups have become much less likely to hold a medium skilled job, and more likely to hold either a low or high skilled job. Job polarization largely reflects changes in worker outcomes, suggesting changes in labour reallocation, rather than changes in the composition of the workforce.[21]

2.3.2 Contribution of Demographic Groups

The previous section established that job polarization reflects changes in worker outcomes at the group-level. The following two sub-sections aim to identify these changes in terms of shifts in employment shares. In this sub-section, I ask how important demographic groups are in accounting for employment share changes.

Tables 2.2, A.2, and A.3 present contributions of each group to the total change, in terms of each decomposition component separately, of the respective employment share. Contributions are computed as the percentage of the group-specific component of the total change in employment share j over the respective period. I focus here on the decline in medium skilled employment, the salient feature of job polarization, showing results for low and high skilled employment in tables A.2 and A.3 in the appendix.

Each table is composed as follows: the upper panel reports results for women,

[21]NESPD and ASHE contain no information on educational attainment. Conducting the decomposition for demographic groups without controlling for education effectively attributes groups' changes in educational attainment to propensity changes. Propensity changes should thus be understood as changes in worker outcomes for meaningful demographic groups not conditional on their educational attainment.

and the lower for men. The two top rows give the total change for the employment share over each period in percentage points. Column one indicates the age group. The contribution term is given in column two, and columns three to 10 give, for each period, the respective contribution as a percentage of the total change in employment share j.[22] I focus on discussing propensity changes as, first, these have been shown in the previous sub-section to be most important for job polarization, and, second, my aim is to identify changes in worker outcomes associated with job polarization. I briefly discuss compositional changes towards the end.

Table 2.2 shows each group's contribution to the decline in the population share for medium skilled employment. Focusing on the overall period 1981 to 2015, we see that all groups contribute to the 12.7 pp decline via propensity changes: the distributions for all demographic groups shift away from medium skilled jobs. The magnitude of contributions varies substantially, however. The shift away from medium skilled employment at the aggregate level is driven primarily by men: men's propensity changes contribute 61.6 percent to the overall decline, while women only contribute 24.2 percent, mostly reflecting changes concentrated among young women. Corresponding shifts for prime aged and older women contribute little to the overall decline in medium skilled jobs.

One may wonder whether these contributions to the overall change hide variation across sub-periods. That is, do employment distributions of young women and men shift away from medium skilled jobs in each sub-period, or does the overall pattern reflect more complex changes? Examining contributions by sub-period we observe that, for men, propensity contributions largely reflect a consistent pattern across periods. This consistent contribution signals a continuing shift of their employment distribution away from medium skilled jobs. Yet, the magnitude of their contribution varies somewhat over time, and is comparatively low in the first period, predating the overall period starting in 1981, and again in the most recent period.

In contrast, women do not exhibit such sustained shifts in the distribution. Young women start contributing more substantially to the decline from the 1990s recession onwards, when they become less likely to be employed in medium skilled jobs. The contribution of prime aged and older women appears more important when considering sub-periods than is implied by contributions for the overall period. This discrepancy reflects two factors: first, women of these age groups experience partly

[22]For instance, the contribution of young women in the first period from 1975 to 1979 in terms of propensity changes to the total change of medium skilled employment of -2.2 percentage points is -0.3. That is, young women counteract the decline by a very small amount, 0.3 percent of 2.2 percentage points, as they become more likely to hold a medium skilled job over this period. Note that expressing contributions in terms of percentage of the total change implies that the sum of contribution over all groups in each period is 100.

offsetting distributional shifts, so their contributions cancel out to some extent. For instance, in the long period from 1992 to 2008 prime aged and older women become more likely to hold a medium skilled job, thereby counteracting the overall decline in this period. Second, the larger contribution also reflects the fact that they comprise an increasing share of the workforce. The increasing weight is not taken into account when considering the overall period. Interestingly, their contribution is particularly large in the most recent period. Evidently, although their overall contribution appears comparatively small, variation across sub-periods reveals a growing importance for distributional shifts experienced by these groups.

We see that the decline in medium skilled employment occurs for the most part because young women, and male workers of all ages, are less likely to hold a medium skilled job. Men contribute most to this decline, because they experience a general and continued shift away from medium skilled employment, while the more limited contribution of women reflects a more varied pattern of initially offsetting distributional shifts, revealing a growing importance in more recent years. The pattern for the overall period should thus not be taken to suggest that prime aged and older women are unimportant for current, and potential future job polarization.

Job polarization implies a shift in the employment distribution away from medium skilled towards low and high skilled jobs. How did groups contribute to the respective increases in low and high skilled employment shares? Tables A.2 and A.3 in the appendix show contributions. Again, I begin by focusing on the overall period 1981 to 2015, during which low and high skilled employment shares increase by 8.8 and 3.9 pp respectively.

Start with contributions to the change in the low skilled employment share. The rise in low skilled employment occurs as employment distributions of men and young women shift towards low skilled jobs, despite offsetting shifts away from low skilled jobs experienced by prime aged and older female workers. In particular, it is the increasing propensity of young workers, of either gender, to hold a low skilled job, which is accounting for most of the increase. Propensity changes for these two groups can account for two thirds of the increase.

Young workers also exhibit sustained distributional shifts towards low skilled employment across sub-periods. The consistent positive and overall large contribution, however, hides some variation in the relative importance of young women and men. Shifts towards low skilled jobs start as a predominantly young and male phenomenon during the 1980s. Over time the relative contribution of young men declines, and young women gain in importance.

The negative contribution of prime aged and older female workers reveals mixed contributions across sub-periods. Interestingly, the sharp increase in low skilled employment following the Great Recession reflects, on the one hand, the growing importance of young women, but also a trend reversal for women aged 31 or older. Prime aged and older women are now becoming more likely to hold a low skilled job.

The contribution of groups to the rise in high skilled jobs contrasts with their contribution to low skilled employment. While the rise in low skilled jobs is mainly a phenomenon afflicting young workers, the rise in high skilled jobs is driven by prime aged and older workers, and especially forcefully by those who are male. Most of the 3.9 pp increase over the period from 1981 to 2015 reflects prime aged and older male workers becoming more likely to hold a high skilled job. Their propensity increase can account for 83.6 percent of the rise in the employment share. Women of the same age groups can only account for 35.8 percent. Young women contribute little via propensity changes, and distribtional shifts for young men even counteract the rise in high skilled employment shares.

The relatively large contributions of prime aged and older workers reflect a largely stable pattern across periods. Important deviations from this pattern occur in the most recent period: The decrease in high skilled employment following the Great Recession is driven by propensity changes for young male and female workers, whose contributions gain in importance, and older female workers, whose distribution starts shifting away from high skilled jobs. The decline occurs despite prime aged and older male workers continuing to become more likely to hold a high skilled job.

These findings suggest that the increasing importance of distributional shifts towards low rather than high skilled jobs after the Great Recession is driven by women. This provides further evidence for the potential importance of distributional shifts, and the corresponding process of labour reallocation, experienced by women for current and future polarization.

An additional, potentially interesting feature relates to the fact that distributional shifts towards low skilled jobs appear to decline with age, and vice versa for shifts towards high skilled jobs. If distributional shifts reflect increasing educational attainment, one would expect young workers becoming more important for shifts towards high skilled jobs over time. Arguably, most workers attain higher education in their early twenties, and so are included among young workers. The fact that young workers are of minor importance for, or even counteract, this shift in early as well as in later periods raises the question whether the increase in low skilled employment might reflect workers becoming more likely to hold low skilled jobs primarily at the beginning of their career, but not throughout. In this case, distributional shifts towards low skilled employment may be accompanied by a rising incidence of occupational upgrading.

To understand the modest contribution of compositional changes, one needs to look at compositional changes across groups in combination with changes in demographic weights.[23] Figure A.3 in the appendix shows changes in the population weights of demographic groups. Propensities in 1981 can be seen in table 2.3, discussed in more detail in the next sub-section.

Overall, the period 1981-2015 witnesses both an aging of the working population and a shift towards female workers. The weight of prime aged and older women increases, while young women remain fairly constant and male workers, especially young workers, see their weights declining. Compositional changes contribute to the decline in medium skilled employment as the composition shifts away from workers with comparatively high propensity for medium skilled jobs, i.e. young workers, towards workers with comparatively low propensity, i.e. middle aged and older women. The shift towards women also contributes to the rise in low skilled jobs, but counteracts the rise in high skilled jobs, as women are comparatively more often employment in a low, and less often in a high skilled job.

[23]Note for the interpretation of group specific contributions via compositional changes that the mechanics of the decomposition imply that a demographic group always contributes to the decline (increase) whenever its share declines (increases). To understand how demographic changes affected employment shares, one needs to sum over group-specific compositional contributions.

Table 2.1: Decomposition of employment share changes

	Period	π_0^j	$\pi_1^j - \pi_0^j$	Decomposition Term: Composition	Decomposition Term: Propensity	Decomposition Term: Interaction
Low Skilled	1975-1979	16.8	1.7	0.7	0.9	0.0
	1979-1981	18.4	0.5	0.1	0.3	0.0
	1981-1990	18.9	0.8	1.1	-0.2	-0.1
	1990-1992	19.6	1.1	0.5	0.5	0.0
	1992-2008	20.7	2.8	1.3	2.3	-0.8
	2008-2009	23.5	0.4	0.1	0.3	0.0
	2009-2015	23.9	3.8	0.4	3.3	0.1
	1981-2015	18.9	8.8	4.0	6.6	-1.8
Medium Skilled	1975-1979	69.5	-2.2	-0.6	-1.6	0.0
	1979-1981	67.3	-1.0	-0.1	-0.9	0.0
	1981-1990	66.3	-1.8	-1.0	-1.0	0.2
	1990-1992	64.6	-2.0	-0.5	-1.4	0.0
	1992-2008	62.6	-4.8	-1.2	-4.2	0.7
	2008-2009	57.8	-0.9	-0.1	-0.8	0.0
	2009-2015	56.9	-3.3	-0.2	-3.1	0.0
	1981-2015	66.3	-12.7	-3.3	-10.9	1.5
High Skilled	1975-1979	13.7	0.6	-0.1	0.7	0.0
	1979-1981	14.3	0.5	0.0	0.5	0.0
	1981-1990	14.8	1.0	-0.1	1.2	0.0
	1990-1992	15.8	0.9	0.0	0.9	0.0
	1992-2008	16.7	2.0	-0.1	1.9	0.1
	2008-2009	18.7	0.5	0.0	0.5	0.0
	2009-2015	19.2	-0.4	-0.2	-0.1	-0.1
	1981-2015	14.8	3.9	-0.7	4.3	0.3

Decomposition results for low, medium, high skilled employment shares over period 1975 to 2015 in percentage points. Decomposition is constructed according to equation 3.5. Column 'Composition' states decomposition term $\sum_{g=1}^{G} \Delta w_{g0} \pi_{g0}^{j}$, column 'Propensity' states decomposition term $\sum_{g=1}^{G} w_{g0} \Delta \pi_{g1}^{j}$, and column 'Interaction' gives decomposition term $\sum_{g=1}^{G} \Delta w_{g1} \Delta \pi_{g1}^{j}$. Six groups are considered: three age groups (18-30, 31-50, 51-65) and two gender groups. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Sub-periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Table 2.2: Group Contributions to Decomposition of Medium Skilled Employment Share

Age	Period	1975-1979	1979-1981	1981-1990	1990-1992	1992-2008	2008-2009	2009-2015	1981-2015
	Total Change	-2.2	-1.0	-1.8	-2.0	-4.8	-0.9	-3.3	-12.7
Age	Contribution of female workers								
18-30	Composition	-21.6	-34.4	-84.8	9.4	31.8	14.0	-5.9	1.4
	Propensity	-0.3	7.3	-5.4	20.0	25.8	24.1	26.8	19.3
	Interaction	0.0	0.3	-1.0	-0.4	-4.3	-0.5	0.8	-0.4
31-50	Composition	-39.0	-21.4	-92.0	-54.2	-30.4	-23.4	18.2	-29.0
	Propensity	30.5	16.4	-21.9	15.8	-7.5	13.5	30.5	2.9
	Interaction	3.0	0.4	-4.0	1.5	-0.9	0.2	-1.3	1.2
51-65	Composition	-10.9	14.2	27.5	-9.3	-34.3	-25.9	-25.4	-20.2
	Propensity	9.4	12.0	6.2	8.4	-3.7	0.8	6.2	2.0
	Interaction	0.5	-0.4	-0.7	0.4	-1.6	0.0	0.9	1.2
Age	Contribution of male workers								
18-30	Composition	24.2	12.4	30.9	51.8	58.4	39.7	-6.8	36.5
	Propensity	6.6	12.5	18.6	9.1	35.9	23.2	5.9	24.0
	Interaction	-0.2	-0.1	-0.7	-0.7	-8.4	-1.1	0.2	-8.0
31-50	Composition	22.6	-36.4	20.9	-2.1	18.8	17.9	34.8	21.8
	Propensity	14.1	26.1	42.5	14.6	24.8	22.3	19.7	24.9
	Interaction	-0.3	0.5	-0.8	0.0	-1.2	-0.2	-1.4	-3.5
51-65	Composition	51.9	79.4	152.2	31.3	-18.2	-12.5	-8.5	15.3
	Propensity	10.6	11.9	16.5	4.6	13.5	7.8	5.3	12.7
	Interaction	-0.9	-0.8	-3.8	-0.3	1.5	0.1	0.2	-2.2

Decomposition contributions to change in medium skilled employment share in percent. Decomposition is constructed according to equation 3.5. Six groups are considered: three age groups (18-30, 31-50, 51-65) and two gender groups. Results show group's composition, propensity, or interaction term as percentage of medium employment share change. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Sub-periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

2.3.3 Change in Propensities of Demographic Groups

The previous sub-sections show that job polarization mainly results from changes in worker outcomes, i.e. shifts in the employment distribution of demographic groups, and they discuss how these shifts contributed to job polarization. This leaves open the question how exactly each group's employment distribution is shifting. The current sub-section provides an answer. The following section will then examine the underlying process of labour reallocation giving rise to these shifts.

Table 2.3 displays propensities in 1981, and propensity changes from 1981 to 1990, 2008 and 2015, for each demographic group. These changes show how population distributions evolve over subsequent expansionary periods. Again, I focus on periods starting in 1981 as the counterfactual analysis examines the corresponding shifts in terms of labour reallocation.

In 1981 medium skilled employment is the most common type of employment for all demographic groups, but substantial differences apply across genders: employment is more clearly concentrated in medium skilled jobs for men than for women, and low skilled jobs are much more common among women than men. Men of all age groups are also more likely than women to have a high skilled job, but the difference is less pronounced.

Considering the change in propensities to 2015, table 2.3 confirms that all groups shift away from medium skilled jobs. This decline is strongest for workers who initially have a high propensity to work in medium skilled jobs, i.e. young women, and men. Table 2.3 confirms prime aged and older women only have a small propensity contribution to the decline in medium skilled jobs because they experience only a small and discontinuous decline in their respective propensity, starting from an already low level. The large contribution of young workers to the rise in low skilled jobs is revealed to reflect very large increases in the propensity to hold a low skilled job. Young workers shift away from medium skilled jobs just as much as they shift towards low skilled jobs. Correspondingly, prime aged and older workers shift mostly towards high skilled jobs.

Comparing changes for the periods ending in 1990, 2008, and 2015 for men, wee see that men exhibit a stable pattern of job polarization. Notable is the large increase in the propensity for low skilled employment for young male workers following the Great Recession. Also, job polarization experienced by men appears to become more pronounced from the 1990s onwards.

Women exhibit a discontinuous pattern of distributional shifts. No pattern of job polarization is apparent for young and prime aged women during the 1980s, and prime aged women only experience a decline in medium skilled employment after the Great Recession ends. Young women experience an even larger acceleration in their shift towards low skilled jobs over the last period.

Changes in employment propensities confirm that all groups experience job polarization over the period from 1981 to 2015. That is, all groups experience a decline in the share of medium skilled employment. Yet, differences obtain with regard to the timing and shifts to low and high skilled jobs. Prime aged and older men shift towards both low and high skilled jobs. Young male workers shift exclusively towards low skilled jobs, away from both medium and high skilled jobs. Young women see a similar increase in their probability to hold a low skilled job. They mainly shift from medium to low skilled jobs, exhibiting a near constant share in high skilled employment. In contrast to young workers, prime aged and older women shift towards high skilled jobs. Also, they experience only a modest decline in the propensity for medium skilled jobs. Although they also shift from low towards high skilled jobs in earlier periods, these shifts are partly offset in more recent years.

2.3.4 Summary of Decomposition Results

This section set out to address two questions: to what extent does job polarization reflect changes in worker outcomes, and which propensity changes are associated with job polarization? Summarizing decomposition results, one can conclude that: (a) job polarization mainly reflects changes in propensities for demographic groups; (b) all groups experience job polarization over the period 1981 to 2015; (c) all groups contribute to job polarization, but changes experienced by male workers and young women appear to be most important, while prime aged and older women become more important in recent years; (d) job polarization reflects shifts in the employment distribution away from medium skilled jobs, with most changes occurring during expansionary periods, and showing no sign of slowing down. The magnitude of the decline in medium skilled employment, and the extent of shifts to either low or high skilled jobs, vary by group. Young workers shift mostly to low skilled jobs, especially in recent years. Prime aged workers shift mostly to high skilled jobs. Recent periods exhibit a change in distributional shifts, witnessing an increasing importance for shifts to low skilled over high skilled jobs.

Overall, we observe that job polarization reflects changes in worker outcomes for all demographic groups, although the changes experienced by workers in each group differ substantially. These differences in distributional shifts are suggestive for differences in the process of labour reallocation. The next section sets out to explain these shifts in terms of labour reallocation.

Table 2.3: Group Employment Share Changes

Age	Status	1981 Level	1981-1990 Change	1981-2008 Change	1981-2015 Change
			Women		
18-30	Low	28.5	0.2	7.2	18.1
	Medium	60.5	0.7	-9.8	-18.6
	High	11.0	-0.9	2.6	0.5
31-50	Low	37.6	-4.3	-6.1	-3.9
	Medium	49.8	2.1	2.2	-2.1
	High	12.5	2.1	3.9	6.0
51-65	Low	42.5	-0.8	-4.6	-0.9
	Medium	48.8	-1.3	-1.3	-3.0
	High	8.8	2.1	5.8	3.9
			Men		
18-30	Low	8.7	2.5	14.0	19.3
	Medium	76.5	-1.8	-13.4	-16.8
	High	14.7	-0.7	-0.6	-2.4
31-50	Low	6.7	0.6	2.7	4.4
	Medium	73.3	-2.8	-8.3	-11.8
	High	20.0	2.2	5.6	7.3
51-65	Low	8.5	0.1	1.1	2.0
	Medium	76.3	-1.9	-8.7	-10.8
	High	15.2	1.9	7.6	8.8

Employment share changes for low, medium, and high skilled employment in percentage points, by demographic groups. Six groups are considered, three age groups (18-30, 31-50, 51-65) and two gender groups. Column 'Level' gives employment share in 1981. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods refer to expansionary periods, starting at through of 1981 recession to peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

2.4 Job Polarization and Labour Reallocation

In the previous section I identified shifts in the employment distributions of demographic groups which account for job polarization at the aggregate level. This serves to demonstrate that polarization reflects changes in worker outcomes, rather than compositional changes, and to identify the related changes in terms of distributional shifts. Ultimately, I want to examine the impact of polarization at the worker-level. Distributional shifts provide only a rough indication of this impact. They constitute a point-in-time measure at the group-level, but they do not reveal how workers are affected over time. The impact is best understood in terms of labour reallocation, an inherently dynamic measure at the worker-level. In this section, I want to examine how these distributional shifts are associated with changes in the process of labour reallocation. That is, I address the question how job polarization relates to labour reallocation.

To link labour reallocation to job polarization I examine worker transition rates. I argue that, in order to understand how workers are affected by job polarization, one needs to examine which changes between labour market states workers experience over time, and I refer to these changes between labour market states as labour reallocation. In particular, I measure labour reallocation as the probability that a worker changes from one state in the current year to another state in the next year. This probability is given by the worker transition rate.[24]

How do worker transition rates relate to propensity changes? I argued above that changes in propensities are only suggestive for the impact at the worker level as the same propensity change is compatible with several changes in transition rates. Take, for instance, propensity changes of prime aged male workers. The previous section demonstrated that at a point-in-time in 2015, compared to 1981, these workers are less likely to hold a medium skilled job and more likely to hold either a low or high skilled job. These changes could reflect workers in medium skilled jobs being more likely to upgrade and get a high skilled job, or workers from low skilled jobs being less likely to upgrade and get a medium skilled job. They could partly reflect prime aged male workers employed in medium skilled jobs entering non-employment more often, and thus experiencing more frequent non-employment spells, or they are less likely to return from non-employment, and thus experience longer non-employment spells. To identify transition rate changes which can account for the observed propensity changes in this context I conduct a counterfactual analysis, described below in more detail.

The basic premise of the counterfactual analysis is as follows. The effect of a transition rate change on the change in employment shares depends on the change itself, the stocks of workers in each labour market state, and all other transition rates. To examine this effect, I express labour market stocks in terms of a Markov process, a function of stocks and transition rates for all labour market states. Using this Markov process I compute and compare benchmark and counterfactual propensity changes. Benchmark propensity changes reflect changes in all transition rates, and counterfactual propensity changes reflect changes in all transition rates except

[24] Note that, as my interest lies in labour reallocation being related to distributional changes in labour market states, I abstract from job to job transitions within the same labour market state. Also, as ASHE and NESPD feature annual observations, I do not observe, and thus abstract from, high frequency transitions occurring between survey dates. Arguably this is especially relevant for transitions to non-employment. I discuss issues related to time aggregation in more detail below.

for one particular rate. This rate is held constant at a level in the baseline period, typically the first expansionary period. The counterfactual propensity change can therefore be thought of as the propensity change which would have occurred if this particular rate's change would not have occurred during expansionary periods. I focus on expansions because recessions are short-lived, and so have a small impact on long-term employment share changes, and because labour reallocation during recessions is likely to be different.[25] The difference to the benchmark change indicates the effect of this rate's change. Transition rate changes are taken to contribute to job polarization if the counterfactual propensity change is small compared to the benchmark change, i.e. if holding this rate constant at the baseline level substantially mitigates the observed propensity change.

I conduct the counterfactual analysis at the demographic group level to examine how transition rate changes contributed to the propensity changes discussed in the preceding section. The previous analysis suggests the following focus on propensity changes. I showed that job polarization is experienced by all groups, and all groups contribute to aggregate job polarization, but both the pattern of distributional changes and the importance for aggregate job polarization differ across groups. Varying patterns in job polarization across demographic groups are suggestive for possible differences in labour reallocation and so warrant examining this link separately for each group. However, some groups warrant particular attention as they are both disproportionately affected and contribute most to job polarization, i.e. male workers in general, and young female workers, especially since the 1990s recession.

I start by presenting the counterfactual analysis in more detail in sub-section 2.4.1. To provide the necessary context for the interpretation of results, I discuss how transition rates changed over time in sub-section 2.4.2. Sub-section 2.4.3 presents and discusses results.

2.4.1 Conducting the Counterfactual Exercise

The aim of the counterfactual exercise is to compare benchmark and counterfactual changes in employment shares. Employment shares are based on the stock of workers

[25] For instance, Haltiwanger [2011] argues for the US that earnings losses associated with displacement are more severe during recessions, and non-employment duration increases, consistent with the view that worker separations during expansions are more likely to be associated with workers reallocating from low-productivity to high-productivity firms. See also Elsby et al. [2011]; Gomes [2012]; Elsby et al. [2015]; Shimer [2012]; Fujita and Ramey [2007] for a discussion of cyclical fluctuations in unemployment in- and outflow rates, and Fallick et al. [2012]; Moscarini and Postel-Vinay [2012]; Fujita and Moscarini [2013] for a discussion of cyclical fluctuations in job-to-job transition rates and worker recall.

in each labour market state.[26] To prepare for the analysis, I compute annual stocks of workers in each employment state as well as the stock of non-employed workers. The stock of non-employed workers serves to account for transitions between employment and non-employment. For technical reasons, I also compute stocks for sample entry and exit.[27] The different labour market states are denoted as follows: Employment in low, medium, or high skilled job (L,M, or H), non-employment (N), and entry and exit (E and X). For $A, B \in \{L, M, H, N, E, X\}$, denote the stock of workers in labour market state A at time t as A_t, the stock of workers in labour market state A at time t and in B at time $t+1$ as A_t^B, and the transition rates from state A to B at time t as λ_t^{AB} as the ratio of the stock of workers in state A at time t and in state B at time $t+1$ to the stock of workers in state A at time t, i.e. $\lambda_t^{AB} = A_t^B / A_t$.

I conduct the analysis at the demographic group level to abstract from compositional changes. I therefore create the above described stocks and transition rates for six sub-samples, comprising observations for each demographic group, and repeat the analysis for each group in turn. To motivate the analysis, it helps to first express future annual stocks as a function of current labour market stocks and transition rates:

$$\begin{bmatrix} E \\ L \\ M \\ H \\ N \\ X \end{bmatrix}_{t+1} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \lambda^{EL} & \lambda^{LL} & \lambda^{ML} & \lambda^{HL} & \lambda^{NL} & 0 \\ \lambda^{EM} & \lambda^{LM} & \lambda^{MM} & \lambda^{HM} & \lambda^{NM} & 0 \\ \lambda^{EH} & \lambda^{LH} & \lambda^{MH} & \lambda^{HH} & \lambda^{NH} & 0 \\ \lambda^{EN} & \lambda^{LN} & \lambda^{MN} & \lambda^{HN} & \lambda^{NN} & 0 \\ \lambda^{EX} & \lambda^{LX} & \lambda^{MX} & \lambda^{HX} & \lambda^{NX} & 0 \end{bmatrix}_t \begin{bmatrix} E \\ L \\ M \\ H \\ N \\ X \end{bmatrix}_t$$

This corresponds to a Markov system, with $\boldsymbol{S}$ denoting the column vector of stocks and $\boldsymbol{\Lambda}$ the transition matrix consisting of annual transition rates:

$$\boldsymbol{S_{t+1}} = \boldsymbol{\Lambda_t S_t}$$

Taking $\boldsymbol{S_0}$ as given, stocks in year T can be written as a function of transition rates as follows:

[26] The counterfactual analysis largely follows Cortes et al. [2014]. A similar exercise is conducted by Shimer [2012]. See also Fujita and Ramey [2007, 2009] for a general discussion.

[27] The entry stock refers to the number of workers entering the sample in the following year. The exit stock refers to the number of workers who left the sample in the previous year. I only compute these stocks to account for transitions in and out of the sample. These stocks are taken as exogenous. Obviously, transition rates into sample entry or out of sample exit are zero.

$$\boldsymbol{S_T} = \boldsymbol{\Lambda_{T-1}\Lambda_{T-2}...\Lambda_t S_0} \tag{2.2}$$

Starting from this description, I conduct the counterfactual analysis as follows. First, I separately compute benchmark stocks, denoted $\boldsymbol{S^{BM}}$, as well as counterfactual stocks for the transition rate from A to B, denoted $\boldsymbol{S^{AB}}$ for $A, B \in \{L, M, H, N, E, X\}$, in year T. Second, I compute employment shares for each labour market state and compare the benchmark and counterfactual change in employment shares over several periods. I discuss the computation of benchmark and counterfactual stocks in turn.

Benchmark Stocks

Benchmark stocks $\boldsymbol{S^{BM}}$ are computed based on equation 2.2, replacing the matrix of annual transition rates with transition rates averaged over the respective period. I use period transition rates because I am interested in long-term changes in labour reallocation, reflecting changes in reallocation across periods. I take it that average transition rates provide a better summary for these changes than annual transition rates.[28] I demonstrate below that employment share changes based on average transition rates provide a good approximation to changes based on annual transition rates.

Suppose period i lasts k years. Denote the matrix with period-average transition rates as $\boldsymbol{\bar{\Lambda}_i}$, such that for each element $\bar{\lambda}_i^{AB} = \left(\sum_{j=1}^{k} \lambda_j^{AB}\right)/k$. In each year of period i the matrix of period-averaged transition rates $\boldsymbol{\bar{\Lambda}_i}$ is

$$\boldsymbol{\bar{\Lambda}_i} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \bar{\lambda}^{EL} & \bar{\lambda}^{LL} & \bar{\lambda}^{ML} & \bar{\lambda}^{HL} & \bar{\lambda}^{NL} & 0 \\ \bar{\lambda}^{EM} & \bar{\lambda}^{LM} & \bar{\lambda}^{MM} & \bar{\lambda}^{HM} & \bar{\lambda}^{NM} & 0 \\ \bar{\lambda}^{EH} & \bar{\lambda}^{LH} & \bar{\lambda}^{MH} & \bar{\lambda}^{HH} & \bar{\lambda}^{NH} & 0 \\ \bar{\lambda}^{EN} & \bar{\lambda}^{LN} & \bar{\lambda}^{MN} & \bar{\lambda}^{HN} & \bar{\lambda}^{NN} & 0 \\ \bar{\lambda}^{EX} & \bar{\lambda}^{LX} & \bar{\lambda}^{MX} & \bar{\lambda}^{HX} & \bar{\lambda}^{NX} & 0 \end{bmatrix}_i$$

Recall that, as described above, I classify the total sample period from 1975 to 2015 into 3 recessionary and 4 expansionary periods. Period 1, 3, 5, and 7 correspond to expansionary periods, and periods 2, 4, and 6 correspond to recessionary periods.

[28] I follow here Cortes et al. [2014], who use period transition rates for the same reason. In line with the decomposition, I choose as periods for the calculation of transition rates: 1975-1978, 1979-1981, 1982-1989, 1990-1992, 1993-2007, 2008-2009, 2010-2015. This implies that for transition rates expansionary periods start the year after the through of the preceding recession, and end the year before the peak of the subsequent recession.

Period 1 serves as the benchmark period. Denote the number of years each period i lasts as y_i. Benchmark stocks after the end of some period l are computed according to equation 2.2, replacing $\boldsymbol{\Lambda_t}$ with $\boldsymbol{\bar{\Lambda}_i}$ for each year of period i:

$$\boldsymbol{S_l^{BM}} = \boldsymbol{\bar{\Lambda}_{l-1}}^{y_{l-1}} \boldsymbol{\bar{\Lambda}_{l-2}}^{y_{l-2}} ... \boldsymbol{\bar{\Lambda}_1}^{y_1} \boldsymbol{S_0} \tag{2.3}$$

Counterfactual Stocks

Counterfactual stocks $\boldsymbol{S^{AB}}$, i.e. stocks associated with the counterfactual transition rate from state A to B, are computed based on equation 2.2, replacing the matrix of annual transition rates with transition rates averaged over the respective period, additionally holding one transition rate at a time constant at its average level in period 1. Denote the matrix of counterfactual transition rates for period i, corresponding to the matrix of transition rates holding the transition rate from labour market state A to B constant, as $\boldsymbol{\Lambda_i^{AB}}$.[29] As an explanatory example, suppose I hold the transition rate from medium to low skilled employment constant, i.e. I compute $\boldsymbol{S^{ML}}$.

The rationale for computing counterfactual stocks is to compute how stocks would have evolved had a particular transition rate remained at its level in period 1 during subsequent expansionary periods.[30] As rates adjust in recessionary periods, this implies I use the same transition rate matrix for recessionary periods for counterfactual and benchmark stocks. Formally, $\boldsymbol{\Lambda_i^{AB}} = \boldsymbol{\bar{\Lambda}_i}$ for $i = 2, 4, 6$.[31]

[29]Note that I treat sample exit rates as exogenous and so do not conduct counterfactual analyses for these rates. Sample exit rates reflect simply reflect aging and do not reflect the behaviour of workers. Transitions to sample exit are included only for technical reasons. When using counterfactual transition rates, the stock of sample exits differs from benchmark stocks as exit rates do not adjust endogenously to changes in stocks. As a result, this will lead to technical differences in net sample entry in the case of counterfactual transition rates compared to the benchmark case.

[30]A slightly different reasoning applies when holding sample entry rates constant for prime aged and older workers. Sample entry for these workers mainly reflects ageing. Sample entry rates in this case mainly reflect the combined effect of transition rate changes affecting workers before they entered the current age group. When considering counterfactual entry rates for prime aged and older workers, I therefore allow all non-entry period transition rates to adjust, while *all* non-entry transition rates in period 5 and 7 are fixed at their period 3 level. This counterfactual scenario shows how propensities would have changed if workers continued to enter in later expansionary periods as they did in period 3.

[31]Changes in transition rates during recessions generally do not have a large effect on employment share changes over the periods considered for comparison, as these recessions are short-lived and, for a constant transition matrix, stocks approach steady state after a small number iterations. In an earlier version, conducted for the full sample, I examined the effect of transition rate changes during recessions, examining how employment shares would have changed had a particular transition rate during a recession not changed relative to period 1, or had a particular recession not occurred but instead rates would have remained at the previous period's level. Differences to the benchmark scenario diminished rapidly as stocks approached the steady state determined by the current expansionary period's rates. See Cortes et al. [2014] who make a similar argument.

For expansionary periods the counterfactual transition rate matrix reflects a single transition rate held constant, here from medium to low skilled employment. Transition rates from states other than medium skilled employment simply reflect the period average, i.e. I compute each element $\bar{\lambda}_i^{AB}$ for $A \neq M$ of $\boldsymbol{\Lambda_i^{ML}}$ as above as $\bar{\lambda}_i^{AB} = \left(\sum_{j=1}^{k} \lambda_j^{AB}\right)/k$. For counterfactual stocks associated with the transition rate from medium to low skilled employment, λ_i^{ML} is held constant at its average level in period 1 in all subsequent expansionary periods. For period $i = 3, 5, 7$ the element $\bar{\lambda}^{ML}$ of $\boldsymbol{\bar{\Lambda}_i^{ML}}$ equals the average transition rate for period 1, i.e. $\bar{\lambda}_i^{ML} = \bar{\lambda}_1^{ML}$.

Holding this transition rate constant requires adjusting other transition rates emanating from medium skilled employment, so that the sum of probabilities of either remaining in medium skilled employment or leaving to other states equals one. For non-diagonal transition rates held constant, I adjust the transition rate to remain in the respective state so that the transition rates out of this state sum to one. For instance, suppose that $\bar{\lambda}_i^{ML} > \bar{\lambda}_1^{ML}$, so the counterfactual transition rate matrix exhibits a smaller probability compared to the benchmark matrix to move to low skilled employment out of medium skilled employment in period i. Denote the difference between the benchmark and counterfactual transition rate as $\delta_i \equiv \bar{\lambda}_i^{ML} - \bar{\lambda}_1^{ML}$. I adjust the transition rate to remain in medium skilled employment in period i by adding δ_i, i.e. $\tilde{\lambda}_i^{MM} = \bar{\lambda}_i^{MM} + \delta_i$.

Intuitively, the counterfactual scenario corresponds to the case in which workers are less likely to leave to low skilled employment and more likely to remain in medium skilled employment.[32] In each year of period $i = 3, 5, 7$ the counterfactual transition rate matrix is:

$$\boldsymbol{\bar{\Lambda}_i^{ML}} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ \bar{\lambda}_i^{EL} & \bar{\lambda}_i^{LL} & \bar{\lambda}_1^{ML} & \bar{\lambda}_i^{HL} & \bar{\lambda}_i^{NL} & 0 \\ \bar{\lambda}_i^{EM} & \bar{\lambda}_i^{LM} & \tilde{\lambda}_i^{MM} & \bar{\lambda}_i^{HM} & \bar{\lambda}_i^{NM} & 0 \\ \bar{\lambda}_i^{EH} & \bar{\lambda}_i^{LH} & \bar{\lambda}_i^{MH} & \bar{\lambda}_i^{HH} & \bar{\lambda}_i^{NH} & 0 \\ \bar{\lambda}_i^{EN} & \bar{\lambda}_i^{LN} & \bar{\lambda}_i^{MN} & \bar{\lambda}_i^{HN} & \bar{\lambda}_i^{NN} & 0 \\ \bar{\lambda}_i^{EX} & \bar{\lambda}_i^{LX} & \bar{\lambda}_i^{MX} & \bar{\lambda}_i^{HX} & \bar{\lambda}_i^{NX} & 0 \end{bmatrix}$$

Finally, the transition rate matrix in period 1 simply equals the matrix of period average transition rates, i.e. $\boldsymbol{\bar{\Lambda}_1^{ML}} = \boldsymbol{\bar{\Lambda}_1}$. Using these counterfactual transition rate matrices, counterfactual stocks after the end of period l are computed according to equation 2.4, replacing $\boldsymbol{\Lambda_t}$ with $\boldsymbol{\Lambda_t^{ML}}$ for each year of period i:

[32] Note that I apply a different adjustment for diagonal transition rates held constant, when I add the difference between the benchmark and counterfactual transition rate proportionally to exit rates to other employment types. Another adjustment also applies to sample entry rates for young workers, for whom entry to other employment types and non-employment is adjusted proportionally.

$$\boldsymbol{S_l^{CF}} = \left(\bar{\boldsymbol{\Lambda}}_{\boldsymbol{l-1}}^{\boldsymbol{ML}}\right)^{y_{l-1}}\left(\bar{\boldsymbol{\Lambda}}_{\boldsymbol{l-2}}^{\boldsymbol{ML}}\right)^{y_{l-2}}...\left(\bar{\boldsymbol{\Lambda}}_{\boldsymbol{1}}^{\boldsymbol{ML}}\right)^{y_1}\boldsymbol{S_0} \tag{2.4}$$

Comparison

Having obtained benchmark and counterfactual labour market stocks, I compute employment shares for state $C \in \{L, M, H\}$ as $s_{Ct} = C_t/P_t$, where $P_t \equiv L_t + M_t + H_t$. Denote benchmark and counterfactual changes from 1981 to the end of period T as $\Delta s_{CT}^{BM} = s_{CT} - s_{C0}$ and $\Delta s_{CT}^{AB} = s_{CT}^{AB} - s_{C0}$. The difference between benchmark and counterfactual changes gives the contribution of a change in the transition rate from state A to B for the change in employment share for state C.

I compare benchmark and counterfactual employment share changes for three periods: 1981 to 1990, 1981 to 2008, and 1981 to 2015. These periods start at the through of the first recessionary period and end at the peak of each subsequent expansionary period.[33] I consider these three different periods to examine how changes in labour reallocation associated with job polarization vary over time. The main interest lies in examining changes for the overall period. However, as some transition rates change non-monotonically, I consider shorter sub-periods with the same starting point to be able to see how changes in different periods are driving deviations from benchmark propensity changes.

For instance, suppose the low skilled employment share increases over the respective period, i.e. $\Delta s_{BM}^{CT} > 0$, and suppose the counterfactual change associated with transition rate λ^{ML} is smaller by x percentage points, i.e. $\Delta s_{BM}^{CT} - \Delta s_{ML}^{CT} = x$. The contribution of the change in the transition rate from medium to low skilled employment can account for x percentage points of the increase in the low skilled employment population share: without this change the low employment share would have been lower by x percentage points.

One important limitation applies to this counterfactual analysis.[34] Equations 2.3 and 2.4 imply that each transition rate's contribution is a (nonlinear) function of other transition rates. The contribution of each rate will therefore differ with the level of these rates. If transition rates are correlated, the effect of the systematic co-movement of transition rates on the contribution will be missed by the above described counterfactual analysis. As a result, contributions do not necessarily sum to the total change in employment shares, and some of the change in employment shares may remain unaccounted for. I take it that, despite this limitation, the coun-

[33]Note that benchmark and counterfactual employment shares are the same in 1981, as the transition matrices deviate only from period 3 onward.

[34]This criticism was first raised by Fujita and Ramey [2007, 2009]. See Shimer [2012] for a response.

terfactual analysis is informative about the importance of transition rate changes to job polarization.

2.4.2 Changes in Transition Rates

To provide the context for results for the counterfactual analysis I first discuss changes in transition rates. As I am interested in long-run changes in worker transitions, I focus on period transition rates. To examine changes in average transition rates I estimate a linear probability model for the probability to exit each state to a particular destination. Specifically, at the worker-level I regress transition dummies – taking on value 1 in year t for each worker in state A in year t who is in state B in year $t+1$ – on period dummies for each of the seven periods. I thus run the following regression, separately for each destination and each demographic group:[35]

$$d_{it}^{A}B = \sum_{j=1}^{7} \delta_j p_{it}^{j} + \varepsilon_{it} \tag{2.5}$$

where d_{it}^{AB} is the indicator variable for transitions from state A to B for worker i from the respective demographic group in year t, p_{it}^{j} is a dummy variable for period $j = 1, 2, ..., 7$, and ε_{it} is an error term. Estimates for δ_j give the mean transition probability for period j, i.e. the period transition probability. For the ease of exposition, I display coefficients and confidence intervals at the 95 percent confidence level for each period in figures A.4 to A.13 in the appendix.

The aim of this analysis is to understand how changes in rates affect employment share changes for different demographic groups. I therefore focus firstly on the change in transition rates over time, and secondly compare changes across groups, neglecting differences in levels for the most part. Also, as most coefficient estimates are significant, in the discussion I assume all estimates are significant unless otherwise stated. For the sake of brevity I focus on rates that prove to be important for interpreting counterfactual results.

Transition Rate Changes for Male Workers

The defining feature of job polarization is the disappearance of medium skilled jobs. A plausible proximate explanation for polarization is that workers reallocate from

[35] I exclude observations in years of discontinuities as well as the last sample year, as no flows can be discerned. The sample for non-employment outflow rates starts in 1976 as, due to the sampling scheme, no non-employed workers are observed for the first year. As linear probability models are known to suffer from heteroskedasticity, I use robust standard errors. Note that I do not include a constant.

medium skilled jobs to other types of employment. I therefore start discussing transitions out of medium skilled employment, shown in figure A.5.

If polarization reflects reallocation from medium skilled to other jobs, one would expect transition rates from medium to low or high skilled jobs to increase over time. This is not confirmed by transition rates: no general increase in the probability to reallocate from medium to low or high skilled jobs can be discerned. While these transition rates are very small from the start of the sample period, the probability to leave medium skilled jobs to other types of employment decreases for all age groups in period 3, although this partly reverses in more recent periods. In the most recent period, young male workers are much more likely to reallocate to a low skilled job, while for prime aged and older workers transition rates to high skilled jobs increase significantly. This conforms with decomposition results, finding that young male workers shift primarily to low, and prime aged and older workers to high skilled jobs. In recent years the shift away from medium to either low or high skilled jobs may reflect direct reallocation out of medium skilled employment.

Does job polarization reflect a diminishing reallocation towards medium skilled jobs? Panel (b) in figures A.4 and A.6 shows transition rates from low and high to medium skilled jobs. For all rates there is a marked drop from period 1 to period 3, and rates remain low throughout the sample period. Reallocation from other job types into medium skilled jobs indeed declines. This implies both a decrease in upgrading for workers in low skilled jobs as well as a decrease in downgrading for high skilled workers.

Apart from reallocation between job types, polarization may also reflect an increase in the non-employment inflow rate, suggestive of an increase in job destruction. Panel (d) in figure A.5 shows the corresponding transition rates. There is at best mixed evidence to support this claim. Workers are not generally more likely to move from medium skilled jobs to non-employment. In period 3 the probability to enter non-employment out of a medium skilled job decreases for all age groups. Prime aged workers experience a partial reversal of this decline in subsequent periods, while young workers do exhibit a higher non-employment inflow rate from medium skilled jobs over time.

While the non-employment inflow rate is suggestive of job destruction, the non-employment outflow rate is generally associated with job creation. Changes in the transition rate from non-employment to medium skilled jobs, shown in panel (b) of figure A.7, provide evidence that job polarization may indeed reflect a decline in the creation of medium skilled jobs. Outflow rates from non-employment decline sharply after period 1, and generally remain low throughout the sample period.

Compatible with decomposition findings of a growing shift of younger workers to low skilled employment, we observe that young male workers are more likely to avoid non-employment by working in a low skilled job in recent years. This suggests that part of the recent rise in low skilled employment may reflect young workers reallocating to low skilled jobs from non-employment. The question arises whether this reflects a change in the worker's employment type. Are young workers generally more likely to work in low skilled jobs, and to return to the same job type after non-employment, or are they more likely to return to low skilled jobs after entering non-employment from either medium or high skilled jobs?

Two caveats apply to transition rates into and out of non-employment. First, recall that non-employment is mismeasured, and can reflect migration or self-employment. Changes in non-employment inflow or outflow rates can therefore reflect changes in the probability to emigrate or become self-employed, or to return from either to PAYE registered employment in the UK, rather than changes in the probability to enter or exit non-employment.

The second caveat only applies to non-employment outflow rates: the sampling scheme suggests that part of the initial decline may reflect compositional changes in the stock of non-employed workers. Composition changes can reflect two factors: first, workers exhibit heterogeneous exit probabilities. Second, the exit probability of workers exhibits duration dependence. Either case implies the share of workers with high exit probability declines at the sample start, as the stock of non-employed workers in early years mostly contains workers who just entered non-employment.[36] This implies a drop in the average exit probability. Using period 1 transition rates from non-employment as baseline rates for the counterfactual analysis may therefore overstate the decline and so the contribution of non-employment outflow rates. To alleviate this concern, I additionally use transition rates which have been corrected for compositional changes when conducting the counterfactual analysis. Corrected transition rates confirm that non-employment outflow rates to

[36]Note that hazard rates estimated in chapter 3 show that the exit probability generally decreases with non-employment duration, which suggests either worker heterogeneity or duration dependence are present in the sample.

medium skilled jobs decline after period 1, although the decline is smaller.[37][38]

Finally, I consider transition rates from sample entry. Transition rate changes are shown in figure A.8.[39] Entry rates for young workers generally reflect the first time workers enter the labour market. Prime aged and older workers primarily enter the sub-sample because of aging, so their entry rates largely reflect propensities prior to sample entry.[40] For young workers the figure reveals a dramatic rise in the probability to enter the labour market in a low skilled job and a matching decrease in the probability to enter in a medium skilled one. While young workers at the start of the sample period predominantly begin their career in a medium skilled job, and only rarely in a low skilled one, entries in low and medium skilled jobs are almost equally likely by the end of the sample period. Interestingly, entry into high skilled jobs also decreases by a substantial amount. This suggests that part of the large shift of young male workers towards low skilled jobs results from workers directly entering into low skilled jobs. The fact that older age groups exhibit less pronounced distributional shifts towards low skilled jobs may be taken to imply that workers eventually reallocate from low to other job types.

[37] Worker heterogeneity and duration dependence both imply that workers with longer durations exhibit a lower exit probability. I therefore compute corrected transition rates for period 1 as a weighted average of duration-specific exit probabilities, using weights observed in period 3. The remaining decline in outflow rates from period 1 to period 3 is entirely driven by decreasing duration-specific exit rates, which are independent of worker heterogeneity and duration-dependence. As the decline abstracts from compositional changes, I take the contribution of correct transition rates as a lower bound. Corrected transition rates suggest exit probabilities in period 1 to medium skilled jobs of 0.1908 (young workers), 0.1543 (prime aged workers), and 0.0881 (old workers) respectively. For exits to low and high skilled employment respectively, these change to 0.0230 and 0.0344 (young workers), 0.0162 and 0.0376 (prime aged workers) and 0.0122 and 0.151 (older workers).

[38] Note that findings in chapter 3 support a decline in non-employment outflow rates independently of these issues: estimating hazard rates for male workers entering non-employment in each period shows that workers entering in later periods generally experience longer spells. Longer non-employment durations are consistent with a decline in the non-employment outflow rate. Chapter 3 also shows that longer non-employment durations are driven by a decline in the outflow rate to medium skilled jobs. Results are also roughly in line with changes in non-employment outflow rates for women.

[39] Note that sample entry to non-employment for young workers reflects workers who enter NESPD or ASHE before the age of 18. These workers enter my sample at age 18 in accordance with my definition of the working population as comprising workers aged 18 to 65 years.

[40] Entry rates for prime aged and older workers in period 1 are the exception: NESPD began in 1975 by sampling PAYE registered employees, entry into non-employment is therefore very low for prime aged and older workers compared to later periods. As one can see in panel (d) of figure A.8, entry rates to non-employment rise very sharply from period 1 to period 3. Because entry rates to all states sum to one, this affects all entry rates for period 1. However, this does not affect results for the counterfactual analysis as I use entry rates from period 3 as baseline rates for these workers.

Transition Rate Changes for Female Workers

Patterns for transition rate changes for women are similar to those for men. Here I discuss only notable differences in transition rates which are important for the counterfactual analysis. Transition rates for women are shown in the appendix in figures A.9 to A.13.

Focusing first on transitions out of medium skilled employment, as for men there is no general tendency to leave medium skilled jobs for low or high skilled ones. Only in recent years do women become more likely to directly reallocate to other job types. Recall that one explanation for the decline in medium skilled jobs is an increase in job destruction. There is no evidence for men that job destruction is driving the decline in medium skilled jobs. For women, there is even evidence to the contrary. The probability to enter non-employment from a medium skilled job declines for all age groups.

Transitions out of non-employment also exhibit a very similar pattern to men, with exit rates from non-employment declining from period 1 to period 3 and remaining largely flat thereafter. One exception applies to young male workers. As seen above, they exhibit a fairly steep increase in the probability to leave for a low skilled job in recent years. This also applies to young women, who, additionally, also see a comparatively large increase in their probability to leave to high skilled jobs. Recall that young women, in contrast to young men, do not shift away from high skilled jobs. Note that I display uncorrected transition rates for period 1 here. Using corrected rates non-employment exit rates are generally lower, but still exhibit a decline from period 1 to period 3 for prime aged and older women.[41]

2.4.3 Comparing Benchmark and Actual Propensity Changes

Recall that I use period transition rates instead of annual transition rates when computing both counterfactual and benchmark propensity changes. What I am ultimately interested in, of course, is how changes in long-term transition rates contributed to changes in actual outcomes.

Here, I show that the benchmark propensity changes, based on average transition rates, capture well changes in actual propensities. Table A.4 in the appendix compares actual and benchmark employment share changes for each demographic group. The table shows propensities in 1981 and changes over periods ending in 1990, 2008 and 2015 based on annual and average transition rates. Columns under

[41]Using corrected rates, period 1 transition rates to low, medium and high skilled employment change to 0.0512, 0.0926, and 0.0165 (young women), 0.0726, 0.0991, and 0.0211 (prime aged women), and 0.0367, 0.0437, and 0.0073 (older women) respectively.

the heading 'Average' are based on period-average transition rates. Columns under the heading 'Actual' are based on annual transition rates. The table shows that employment share changes are generally quite similar, whether they are based on annual or average rates, for the period 1981 to 1990, and for 1981 to 2015. Discrepancies are somewhat larger for the period ending 2008 for male workers.[42] While results for the period from 1981 to 2008 should be interpreted with caution, table A.4 by and large confirms that benchmark changes provide a good approximation to actual employment share changes.

2.4.4 Results for Basic Counterfactual Analysis

In this sub-section I discuss results for the basic counterfactual analysis. I focus on male and younger female workers, as these have been shown to be most important for job polarization. Results for male workers and young female workers are shown in tables 2.4 to 2.7. Results for prime aged and older women are shown in the appendix in tables A.5 and A.6.

Each table reports results for counterfactual transition rates for a specific demographic group. Also shown in the table are benchmark propensity changes for the group's respective employment shares. Each panel in the table presents results for a particular employment share, and results are presented separately for each period. The counterfactual transition rate associated with the respective result is shown in column one, which indicates the state out of which transitions are considered, and columns two to thirteen, showing destination states separately for each period. Note that non-employment outflow rates corrected for compositional changes, reflecting the sample start problems discussed above, are denoted as 'non (cr)'. Results are presented as the contribution to the benchmark propensity change in percentage points.[43] For instance, looking at the first panel for low skilled employment of table 2.4, the first column for period 1981 to 1990 in the first row of results shows that the counterfactual transition rate from medium to low skilled employment contributes 1.3 pp to the rise in low skilled employment share of 2.3 pp. Equivalently, results in the second row of the same column show that the change in the transition rate from medium to low skilled employment counteracts the increase by 0.4 pp.

The reasoning of the counterfactual analysis implies that transition rate

[42] These larger discrepancies likely reflect comparatively large jumps in annual transition rates directly at the end of the previous expansionary period. These changes are not due to any known discontinuity.

[43] Note that a positive contribution to a negative benchmark change implies that the transition rate change can account for (part of) that change, whereas a negative contribution to a negative benchmark implies the transition rate change counteracts the benchmark change.

changes are important for job polarization if they have a large contribution to the observed shifts in the employment distribution. As the defining feature of polarization is the decline in medium skilled employment, I focus on transition rate changes contributing importantly to this decline.

Young Male Workers

The decomposition conducted above shows that young male workers shifted from medium and high skilled jobs to low skilled employment. The benchmark scenario implies a decline of almost 18 pp in medium skilled employment share over the period 1981 to 2015, accompanied by a 2.5 pp decline in high skilled employment. Both shifts acted to raise the employment share of low skilled jobs by 20.4 pp.

How are these shifts in the distribution brought about? Table 2.4 shows results for the counterfactual analysis for young male workers.

Young male workers largely shifted from high and medium skilled to low skilled jobs directly at the beginning of their working career. The rise in the transition rates from sample entry to low skilled jobs, and simultaneous decline in entry rates to medium skilled jobs, shown in figure A.8, can account for 10.5 and 9.2 pp of the decline in medium skilled employment.[44] The rise in low skilled employment is additionally fuelled by lower entry rates to high skilled jobs. This rate can account for the entire decline in the high skilled employment share.

This implies that a large part of the decline in medium skilled employment reflects male workers experiencing a worsening of their employment outlook right from the beginning, as they more often started their career in low skilled rather than medium skilled jobs. As these results are unconditional on education, it stands to reason whether this rise falls disproportionately on workers with low education.

Overall, table 2.4 suggests only a limited contribution of reallocation between employment types to the observed shifts. It appears that distributional shifts are not associated with young male workers moving from one job type to another.

There is no evidence that workers reallocate to low or high skilled jobs directly out of medium skilled employment. This also implies that the surge in workers starting their career in a low skilled job is not accompanied by a rising reallocation from low to medium skilled jobs. The respective transition rate contributes 1.4 pp to the shift from medium to low skilled employment. Qualitatively similar results apply for reallocation from low to high skilled jobs. The counterfactual analysis

[44] Recall that counterfactual transition rates require adjusting other transition rates to assure that the sum of rates from sample entry is one. Counterfactual sample entry rates for young workers are adjusted by adding the difference between the benchmark and counterfactual entry rate proportionally to entry rates to other employment types and non-employment.

implies that such transitions do not affect distributional shifts. Addressing the question whether young workers entering in a low skilled job in fact remained in low skilled employment permanently cannot be answered by this analysis, as this requires following workers over longer periods. Examining this question in future research, however, is important for assessing the impact of job polarization on young workers. While there is no evidence of workers being more likely to reallocate directly from medium skilled jobs to low or high skilled ones, there is a smaller impact on distributional changes due to the fact that they are less likely to reallocate out of low or high employment towards medium skilled jobs. These rates contribute 2.1 pp to the shift from medium to low skilled employment, and they imply a shift of 0.8 pp from medium to high skilled jobs.

Is there any evidence that job polarization is driven by job destruction? That is, are young workers experiencing job polarization because they are more likely to loose their medium skilled job? The counterfactual analysis only suggests a minor impact for the transition rate from medium skilled jobs to non-employment. It contributes 1.8 pp to the decline in medium skilled employment. As this reduces the stock of workers holding a medium skilled job while affecting workers in low and high skilled jobs only indirectly, this acts to raise employment shares for low and high skilled jobs somewhat.

The decline in the non-employment outflow rates exhibit more substantial contributions.[45] Examining the contribution of corrected and uncorrected non-employment outflow rates suggests a sizable impact on job polarization due to changes in non-employment outflow rates to low and medium skilled jobs. This is consistent with a decline in job creation for medium skilled jobs. It appears that young male workers shifted from medium to low skilled jobs partly because they are more likely to leave non-employment to a low rather than a medium skilled job. The contribution for the decline in medium skilled employment ranges from 3 to 7.8 pp of the 17.9 pp decline. The question arises whether this reflects workers returning to their previous employment type, or workers settling for low skilled employment after being unable to find a better job. I will address this question below by conducting the counterfactual analysis conditional on previous employment.

The main features are fairly constant over time. The most important contribution stems from direct entry to low rather medium or high skilled jobs over all

[45]Recall that the drop in observed outflow rates from period 1 to period 3 is likely to be upward biased because of compositional changes reflecting either worker heterogeneity or duration dependence. I therefore also use transition rates corrected for compositional changes, for which the decline is entirely driven by changes in duration-specific exit rates. The contribution of non-corrected outflow rates is likely to be overstated, while I take the contribution of corrected outflow rates as a lower bound.

periods. The contribution of non-employment outflow rates is also consistent across periods. The outflow rate to low skilled jobs rises drastically after the Great Recession, and this is reflected in a large rise in the contribution for the overall period relative to the period ending in 2008. Reallocation from medium skilled employment appears relatively unimportant throughout, but the contribution due to the downgrading of workers moving from medium to low skilled jobs is a recent phenomenon. In line with transition rate changes shown in A.5, it contributes little to, or even counteracts, the decline in medium skilled jobs until the most recent period.

Prime Aged Male Workers

In line with decomposition results, benchmark changes in employment share demonstrate a shift from medium to low and mostly high skilled jobs. The benchmark decline in medium skilled employment is 11.5 pp, while the rise in low and high skilled shares is 4.4 and 7.0 pp respectively.

Results for the counterfactual analysis, shown in table 2.5, suggest that overall these shifts largley result from workers being less likely to hold a medium skilled job and more likely to be non-employed as they spend longer time in non-employment. Together with changes occuring in earlier age groups this change can also account for the rise in low and high skilled employment.

Transitions between employment are again of minor importance. As for young male workers, there is no evidence for job polarization reflecting workers moving from medium directly to low or high skilled jobs. Instead, a decrease in the reallocation from other jobs to medium skilled jobs contributes modestly. This appears to be more important for workers transitioning from high skilled jobs. The raste from low and high skilled to medium skilled employment contribute 0.8 and 2.9 pp respectively to the decline in medium skilled jobs.

Job destruction does not play a visible role in the disappearance of medium skilled jobs. The transition rate from medium skilled employment to non-employment exhibits a small and negative contribution. The decline is associated with decreasing outflow rates to medium skilled jobs, however, suggesting a role for job creation. The contribution of this outflow rate ranges from 8.5 to 16.1 pp of the 11.5 pp decline. Even when considering the corrected rate, the drop in this rate can account for almost the entire decline in the medium skilled employment share. This also suggests that some of the rise in low and high skilled employment is purely mechanical: their share rises as the absolute number of workers in medium skilled jobs declines.

Note that these results apply to changes experienced by male workers while they belong to the group of prime aged workers. Transition rate changes may have additional effects as they affect the labour market states of young workers at the time they enter the prime aged group. The combined effects will be captured by entry rates for subsequent age groups, of course. Entry plays only a modest role. It can contribute 2.1 pp to the decline in medium skilled jobs, but 1.7 pp, a relatively large share, to the rise in low skilled employment.[46] This is consistent with the large shifts of young workers towards low skilled employment.

The main features of this pattern also obtain for earlier periods. The decline in medium skilled employment shares is generally accounted for by the decline in probability to exit non-employment by entering a medium skilled job. Also, the decline in transition rates from low and high skilled employment to medium skilled jobs leads to similar results across periods, and non-employment inflow rates are generally unimportant for polarization. There is some additional evidence that medium skilled workers are less likely to upgrade to high skilled jobs in the first two periods, and that this contributes to the modest decline in the high skilled employment propensity.

Older Male Workers

Older male workers experience distributional shifts from medium to low and high skilled jobs, with shifts to high skilled jobs being most important. Over the period from 1981 to 2015, benchmark changes demonstrate a decline in the medium skilled employment share of 11 pp, and increasing shares for low and high skilled employment of 1.8 and 9.2 pp.

Results for the counterfactual analysis are shown in table 2.6. Overall, results suggest the pattern of transition rate change contributions is similar to prime aged male workers. The shift from medium skilled jobs to non-employment largely results because workers, instead of taking on employment in medium skilled jobs, remain longer in non-employment.

The transition rate from non-employment to medium skilled jobs can account for 9.2 to 18.2 pp of the 11 pp decline in medium skilled employment share. There is no evidence that medium skilled employment declines because workers are more likely to loose their medium skilled job. The non-employment inflow rate from

[46] Recall that the contribution of entry transition rates indicates the combined effect of transition rate changes affecting workers before they enter subsequent age groups. Entry rates are held constant at their levels in period 3. Results for the counterfactual are only available for periods ending in 2008 or 2015. Restricting the direct effects of transition rate changes to each demographic group results from the decision to conduct the analysis at the demographic group level. The advantage of distinguishing effects by group is met by the disadvantage of being unable to identify cumulative effects. As an additional analysis I therefore conduct below the counterfactual analysis on the full sample.

medium skilled jobs in fact counteracts the decline by 2.4 pp.

Again, job reallocation between employment types is only of minor importance. There is no evidence that polarization reflects workers leaving medium for low or high skilled jobs. Instead, diminishing reallocation towards medium skilled jobs contributes a modest amount. Changes in transition rates from low or high to medium skilled employment account for 0.7 and 2.4 pp of the 11 pp decline.

The contribution from entry is larger than for younger age groups. This is to be expected as entry reflects the cumulative effect of transition rate changes on workers at younger ages. It can account for 3.7 pp of the decline in medium skilled employment, and some sizeable parts of the rise in low and high skilled employment. On the other hand, this means that counterfactual results suggest that older workers experience most of the propensity changes towards the end of their career.

These main features hold throughout the sample period. Earlier periods exhibit some additional patterns, however, affecting largely changes in low skilled employment shares. The decrease in the propensity for low skilled jobs over the period ending in 1990 is, to a limited extent, due to workers being less likely to move from a medium skilled job to a low skilled one. However, the effect for the overall period is small, as this pattern is partly reversed in the most recent period, when this probability increases. In any case, the effect of this change is small in every period.

Young Female Workers

Young female workers experience a shift from medium mostly towards low skilled employment. Their share of high skilled employment remains almost unchanged. The corresponding benchmark changes show a decline of 19 pp in the employment share for medium skilled jobs, a rise in the low skilled employment share of 18.7 pp, and a very small increase in the high skilled employment share of 0.3 pp.

Results are similar to young male workers. The rise in low skilled employment for the most part reflects young female workers entering the labour market directly in a low skilled rather than a medium skilled job. The corresponding transition rate changes can account for 9.4 or 9.6 pp of the decline in medium skilled employment, and 8.6 pp of the rise in low skilled employment. The contribution of the rising entry rate to low skilled jobs additionally reflects the fact that young women are also less likely to enter in a high skilled job. This rate accounts for 11.7 pp of the 18.7 pp rise in low skilled employment share.

Does reallocation between job types play an important role for young female workers? Young women are more likely to move to low or high skilled jobs directly from medium skilled employment, but this only adds 0.7 and 0.4 pp respectively to the decline in medium skilled employment, and similar amounts to the rise in low and high skilled employment shares. Similar contributions also arise from diminished reallocation from other job types into medium skilled jobs. Overall, direct reallocation across job types does not seem to be important for the distributional shifts experienced by young women.

Young women differ from all other groups experiencing large declines in medium skilled employment in one important regard: their shift from medium skilled jobs does not reflect reallocation towards non-employment. Non-employment outflow rates to medium skilled jobs in fact counteract the decline by 2pp, when considering corrected rates. Uncorrected rates suggest a modest contribution of 4 pp, but, given the likely sample start problems, evidence is not conclusive for a substantial contribution. Outflow rates instead contribute to the rise in low and high skilled jobs. Respective contributions are 2.4 to 5.3 and 0.8 to 2.7 pp. Neither are young women more likely to be non-employed because of an increase in job destruction. Non-employment inflow rates from medium skilled jobs again counteract the decline in medium skilled jobs by 2.1 pp, and, more generally, inflow rates from low and high skilled jobs also contribute to the rise in the respective employment shares.

These main features are again largely consistent across periods. The non-employment outflow rate to medium skilled jobs never contributes substantially to the decline in medium skilled jobs. Most shifts can be accounted for by entry to low rather than medium skilled jobs.

Women shift away from medium skilled jobs not because they are less likely to work. In fact, counterfactual results suggest young women reallocate to low skilled jobs instead, and to lesser extent to high skilled jobs, as they directly enter into low skilled jobs, and leave non-employment more quickly to low and high skilled jobs and even to medium skilled jobs.

Prime Aged and Older Female Workers

Results for women are shown in the appendix in tables A.5 and A.6. As these workers groups only exhibit minor contributions to employment share changes over the period 1981 to 2015, I only comment on major deviations from previously observed patterns. Recall, however, that these groups become more important in recent years.

Both worker groups exhibit modest shifts towards high skilled employment at the expense of low and high skilled employment shares. The most important deviations to prime aged and older male workers are thus the modest decline in medium skilled jobs, and the decrease in low skilled jobs. Which transition rate changes can explain these differential shifts?

The decline in medium skilled employment can be more than accounted for by a decline in the non-employment outflow rate to medium skilled jobs. The magnitude of the corresponding contributions, however, is comparatively small with 5.3 pp in each case. This decline is partly offset by the decline in the non-employment inflow rate from medium skilled jobs. Women are less likely to leave their medium skilled job to non-employment. This rate counteracts the decline by 3.7 pp for prime aged and 5.4 pp for older female workers. As a result, more workers hold medium skilled jobs, which counteracts the rise in high skilled employment shares and contributes to the decline in low skilled employment shares. Additionally, prime aged and older women are less likely to leave non-employment to a low skilled job, which further contributes to the shift from low skilled employment.

One additional noteworthy deviation is that prime aged women do not experience job polarization in earlier periods. Instead, medium skilled employment shares increase. What explains the rise in medium skilled employment? And which changes in transition rates can account for the decline in recent years?

The increases of 0.7 and 2.2 pp over the earlier two periods occur despite prime aged women being less likely to return from non-employment to a medium skilled job. These contributions imply, if anything, a decline in medium skilled employment. Larger contributions adding to the increase accrue to changes in transition rates from medium skilled employment to non-employment and from non-employment to low skilled jobs. Prime aged women are less likely to leave their medium skilled job, and they are less likely to leave non-employment to a low skilled job, and to lesser extent also to a high skilled job. As they spend more time non-employed instead of holding a low or high skilled job, the share of medium skilled employment rises mechanically. As non-employment outflow rates to low and high skilled jobs rise in the most recent period, their contribution to push up the medium skilled employment share diminishes in recent years. This implies that medium skilled employment declines for the overall period because, in the most recent period, prime aged women become relatively more likely again to leave non-employment to low and high skilled jobs.

2.4.5 Results for Counterfactual Analysis using Non-Employment Outflow Rates Conditional on Previous Job Type

The previous analysis shows that one of the most important transition rate changes accounting for job polarization is the decline in the outflow rate from non-employment to medium skilled jobs. This implies that, at any point in time, workers are less likely to be employed in a medium skilled job, and more likely to be non-employed, as they remain in non-employment for longer. It is suggestive to associate this non-employment outflow rate with job creation for medium skilled jobs: job polarization occurs partly as fewer medium skilled jobs are being created, and workers evidently do not compensate for the decline in medium skilled job offers by reallocating to low or high skilled jobs. Does this contribution reflect workers who lost their medium skilled job being less likely to return to a similar job? Or does this contribution reflect changes experienced by workers regardless of their previous job type?

To examine this question, I repeat the counterfactual analysis for conditional exit rates from non-employment. A conditional exit rate is the transition rate from non-employment to other labour market states conditional on the job type they hold before entering non-employment. For simplicity, I refer to a worker whose last job was medium (low, high) skilled as a medium (low, high) skilled worker. This counterfactual analysis allows examining how much of the shift from medium skilled employment to non-employment can be accounted for by medium skilled workers being less likely to return a medium skilled job.

Results are reported in the appendix in tables A.7 for women and A.8 for men. As my interest lies in the contribution of the outflow rate from non-employment to the decline in medium skilled employment shares, I only report results for this employment share. Again, I report results for corrected and uncorrected exit rates. Corrected exit rates are constructed as before.

Results for male workers in table A.8 demonstrate that the contribution of the outflow rate to medium skilled jobs in fact mostly reflects medium skilled workers being less likely to return from non-employment to medium skilled jobs. The sum of conditional contributions across all job types is generally close to the unconditional contributions discussed above. Focusing on the period 1981 to 2015, the contribution of the outflow rate to medium skilled jobs is largest for medium skilled workers of all age groups. The respective contribution is 1.4 to 4.7 pp (3 to 7.88pp for unconditional outflow rate to medium skilled jobs) for young male medium skilled workers, 6 to 11.7 (8.5 to 16.1 pp) for prime aged workers, and 7.2 to 15.3 pp (9.2 to 18.2 pp) for older male workers. This pattern is stable over sub-periods.

Additionally, however, the overall contribution of the unconditional exit rate partly reflects changes for workers coming from low and high skilled jobs, who are also less likely to exit to medium skilled jobs. Relative contributions are especially large for young low skilled workers and older high skilled workers. This makes intuitive sense as the decomposition demonstrated these age groups shift predominantly to such job types. It also demonstrates that reallocation from low or high to medium skilled jobs involving a period of non-employment is an important margin for the decline in medium skilled jobs. Young, prime aged, and older low skilled workers add 0.6, 1.1 and 0.7 pp to the contribution. Their high skilled counterparts add another 0.1, 0.8, and 1.6 pp. Again these relative contributions are largely stable over time.

Table A.7 reports results for women. Results for prime aged and older women are in line with their male counterparts. Recall that for young women outflow rates to medium skilled jobs counteract the decline in medium skilled job, if using corrected transition rates. This implies that young women are more likely to leave non-employment for a medium skilled job. Results demonstrate that this applies to workers regardless of their previous job type, with similar magnitudes for low and medium skilled workers (-1.8 pp for medium skilled workers and -1.4 pp for low skilled workers over the period 1981 to 2015). That is, also low and to lesser extent high skilled workers counteract the decline in medium skilled employment as they are more likely to reallocate following a period of non-employment.

Overall, the shift from medium skilled jobs to non-employment, accounting for a large part of job polarization, mostly reflects medium skilled workers being less likely to return to medium skilled jobs. Generally speaking, however, workers of all job types are affected.

2.4.6 Aggregating Results of Counterfactual Analysis

Previous sections discussed the contribution of transition rate changes to job polarization experienced at the group-level. I conducted the analysis at the group-level to abstract from compositional changes, and to examine variation in labour reallocation changes across workers groups. As job polarization is typically discussed as an aggregate phenomenon, it is instructive to know how much changes at the group-level contribute to aggregate job polarization.

To do so, I aggregate group-specific counterfactual and benchmark employment share changes using annual population weights. This allows comparing aggregate contributions for each transition rate change to aggregate employment share changes.[47] Results are shown in table A.9 in the appendix.

Results at the group-level suggest that the most important transition rate changes accounting for job polarization are the decline in the non-employment out-

[47] Note that these contributions reflect transition rate changes at the group-level. Compositional changes are differenced out, as these are reflected in both the counterfactual and benchmark employment share changes.

flow rate to medium skilled jobs and the relative increase in sample entry rates to low versus medium skilled jobs. The importance of these changes can also be seen at the aggregate level. The decline in non-employment outflow rates can account for 5.3 to 13.8 pp of the aggregate 12.9 pp decline in medium skilled jobs. Sample entry rates to low and medium skilled jobs add around 3.6 pp to this decline, and a similar amount to the rise in low skilled employment.

Reallocation between job types is of minor importance. There is no evidence that polarization occurs as workers directly move from medium to low or high skilled jobs, while there is some evidence that polarization reflects workers being less likely to directly reallocate towards medium skilled jobs. This is again confirmed at the aggregate level. Transition rate changes from medium to low or high skilled jobs only contribute 0.1 pp each to the decline in medium skilled jobs. Transition rate changes from low and high to medium skilled jobs, on the other hand, contribute 0.7 and 1.3 pp. They also contribute similar amounts to the rise in low and high skilled jobs respectively.

There is no evidence that medium skilled jobs disappeared because workers are more likely to leave their medium skilled job and become non-employed, suggesting that job destruction played no important role. The corresponding aggregate contribution is -2.2 pp.

Overall, these results demonstrate that the most important contributions to job polarization, able to account for the bulk of changes in employment shares, are the decline in the non-employment outflow rate to medium skilled jobs experienced by almost all demographic groups, and entry to low rather than medium skilled jobs for young workers.

2.4.7 Results for Counterfactual Analysis at the Aggregate Level

Conducting the analysis at the group-level necessitates dividing the full sample into group-specific sub-samples. For instance, the sample used to conduct the counterfactual analysis for prime aged workers contains all observations for workers aged 31 to 50 years old. In this case, sample entry or exit largely reflects aging of workers. Suppose that workers currently aged 31 to 50 years were greatly affected by a change in the transition rate from medium to low skilled employment when they were young, acting to shift the employment distribution towards low skilled jobs for workers who are about to turn 31. The contribution of this transition rate will not reflect this change in entry rates for prime aged workers in the counterfactual analysis, as entry rates do not adjust to the shift in the employment distribution. Instead, I capture the combined effect of transition rate changes by holding all sample entry rates constant at their level in period 3.

One may be concerned that contributions of transition rates in the main analysis are of limited information, due to omitting those parts of effects accumulating over age groups. To address this concern, I conduct the counterfactual analysis at the aggregate level. That is, I repeat the analysis using the full sample at once, instead of dividing the full sample into sub-samples, and conducting the analysis for each sub-sample separately.

Results are shown in table A.10 in the appendix. It is constructive to compare results to table A.9, which shows results aggregated from the group-level. Conclusions generally remain unchanged when conducting the analysis at the aggregate level compared to the group-level. The non-employment outflow rate to medium skilled jobs remains the most important single transition rate change accounting for job polarization. Reallocation between job types remains of modest importance, and there continues to be no evidence that job destruction of medium skilled jobs is important for polarization.

2.4.8 Summary of Results for Counterfactual Analysis

Overall, the counterfactual analysis suggests that the single most important change in transition rates accounting for job polarization is the decline in the non-employment outflow rate to medium skilled jobs. This implies that job polarization to large parts reflects workers being less likely to hold a medium skilled job, and more likely to be non-employed, as workers remain non-employed for longer, possibly driven by a decline in job creation for medium skilled jobs.

This decline can account for about half to the entire decline in medium skilled employment for most demographic groups, and for similar amounts of the decline at the aggregate level, reflecting mostly changes experienced by male workers, especially prime aged and older ones. The contribution reflects changes experienced by workers of all job types, although medium skilled workers by far exhibit the largest contribution. It is mostly medium skilled workers, who are less likely to return to a medium skilled job, who contribute to the disappearance of medium skilled jobs.

Another important change relates to labour market entry. Young workers, who enter the labour market for the first time, are much more likely to enter in a low rather than a medium skilled job. This change can account for about one fourth of the decline in medium skilled employment. The question arises whether young workers remain permanently in low skilled jobs.

Surprisingly, there is no substantial evidence for job polarization being associated with workers directly reallocating out of medium skilled jobs to other types of employment, although there is some evidence that employed workers of all demographic groups are less likely to directly reallocate into medium skilled jobs. There is also no evidence that job destruction is driving job polarization.

Table 2.4: Counterfactual Results for Young Male Workers

Period	1981-1990				1981-2008				1981-2015			
	Low Skilled Employment											
BM change:	2.3				10.8				20.4			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	1.3	1.0	0.1	0.1	0.7	2.0	0.3	-1.8	0.7	2.1	0.5	-2.1
Med	-0.4	-1.2	0.0	-0.7	0.2	0.9	0.0	0.1	1.4	2.8	0.0	1.3
High	0.0	0.0	-0.2	-0.1	-0.1	0.0	0.0	-0.1	-0.1	0.0	0.2	0.0
Non	-1.1	2.1	0.5	0.2	-0.3	4.4	1.0	1.0	2.3	5.4	1.2	1.0
Non (cr)	0.3	0.6	0.1	0.1	1.2	2.0	0.2	0.3	3.6	2.1	0.1	0.1
Entry	1.9	0.7	0.2	0.0	8.0	4.3	0.6	-0.1	12.5	7.5	1.0	0.0
	Medium Skilled Employment											
BM change:	-1.1				-10.1				-17.9			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	1.1	1.0	0.0	0.1	0.5	2.0	-0.1	-1.5	0.6	2.1	-0.1	-1.7
Med	-0.4	-2.9	-0.9	-1.6	0.2	1.7	-0.5	0.3	1.4	4.2	-0.4	1.8
High	0.0	0.7	1.4	1.0	0.0	1.0	0.1	0.3	0.0	0.8	-0.6	-0.1
Non	-0.9	4.6	-3.6	0.0	-0.2	7.6	-3.7	0.5	1.8	7.8	-2.5	0.6
Non (cr)	0.3	1.4	-0.8	0.0	1.0	3.4	-0.8	0.2	2.9	3.0	-0.3	0.1
Entry	1.7	1.4	-1.0	0.0	6.7	6.1	-1.4	0.0	10.5	9.2	-1.5	0.0
	High Skilled Employment											
BM change:	-1.2				-0.7				-2.5			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	0.2	0.0	0.1	0.0	0.1	0.0	0.4	-0.3	0.1	0.0	0.6	-0.4
Med	0.0	1.7	0.9	0.9	0.0	-0.8	0.5	-0.1	0.0	-1.4	0.4	-0.6
High	0.0	-0.7	-1.6	-1.2	-0.1	-1.0	-0.1	-0.4	-0.1	-0.8	0.8	0.1
Non	-0.2	-2.5	4.2	0.1	-0.1	-3.2	4.7	0.5	0.5	-2.4	3.6	0.4
Non (cr)	0.0	-0.7	0.9	0.1	0.2	-1.5	1.0	0.2	0.7	-0.9	0.4	0.1
Entry	0.3	-0.6	1.2	0.0	1.2	-1.7	2.0	0.0	2.0	-1.6	2.5	0.0

Contributions to employment share changes for young male workers based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to particular transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A, B \in$ (Low, Med, High, Non, Entry). Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Young male workers comprise workers aged 18-30. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table 2.5: Counterfactual Results for Prime Aged Male Workers

Period	1981-1990				1981-2008				1981-2015			
					Low Skilled Employment							
BM change:	0.7				1.8				4.4			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	1.5	0.9	0.2	0.5	1.6	1.0	0.3	0.2	0.7	0.8	0.4	-0.8
Med	-0.8	-1.5	0.0	-0.9	-0.7	-1.1	0.0	-0.6	-0.1	-0.2	0.0	-0.1
High	-0.1	0.0	-0.5	-0.4	-0.3	0.0	-0.6	-0.4	-0.3	-0.1	-0.5	-0.2
Non	-3.2	2.3	0.9	-0.1	-6.2	3.7	1.9	0.1	-4.2	4.5	2.0	0.2
Non (cr)	-0.2	0.7	0.2	0.0	-1.5	2.0	0.5	0.0	-0.1	2.4	0.3	0.1
Entry	n.a.				0.5				1.7			
					Medium Skilled Employment							
BM change:	-2.3				-6.2				-11.5			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	1.2	0.9	0.0	0.4	1.2	1.0	-0.1	0.1	0.5	0.8	-0.1	-0.5
Med	-0.9	-5.9	-1.4	-3.6	-0.7	-4.3	-1.2	-2.4	-0.1	-0.7	-0.2	-0.4
High	0.0	2.1	5.2	3.5	0.1	3.3	5.2	3.0	0.1	2.9	2.8	0.9
Non	-2.5	8.7	-8.6	0.1	-4.4	14.6	-14.5	0.1	-2.7	16.1	-10.5	0.0
Non (cr)	-0.2	2.7	-1.5	0.0	-1.1	8.0	-4.1	0.0	0.0	8.5	-1.7	0.0
Entry	n.a.				0.8				2.1			
					High Skilled Employment							
BM change:	1.6				4.4				7.0			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-0.4	0.0	-0.2	-0.1	-0.4	0.0	-0.4	-0.1	-0.2	0.0	-0.5	0.3
Med	0.0	-4.4	-1.4	-2.6	0.0	-3.3	-1.2	-1.8	0.0	-0.6	-0.3	-0.3
High	0.1	2.1	5.7	3.9	0.3	3.3	5.8	3.4	0.4	3.0	3.3	1.1
Non	0.7	6.4	-9.5	0.2	1.8	10.9	-16.3	0.0	1.5	11.6	-12.5	-0.3
Non (cr)	0.1	2.0	-1.6	0.1	0.4	6.0	-4.7	0.0	0.0	6.2	-2.0	-0.1
Entry	n.a.				0.3				0.4			

Contributions to employment share changes for prime aged male workers based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A, B \in$ (Low, Med, High, Non, Entry). Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Prime aged male workers comprise workers aged 31-50. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table 2.6: Counterfactual Results for Older Male Workers

Period	1981-1990				1981-2008				1981-2015			
					Low Skilled Employment							
BM change:	-0.2				0.7				1.8			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-1.9	-1.0	-0.1	-1.0	2.0	1.1	0.2	1.2	1.5	0.7	0.2	1.1
Med	1.4	1.3	0.0	0.8	-1.0	-1.2	0.0	-0.9	-0.4	-0.7	0.0	-0.8
High	0.3	0.0	0.3	0.2	-0.4	0.0	-0.5	-0.3	-0.5	0.0	-0.4	-0.2
Non	8.7	-3.6	-1.1	0.7	-13.6	5.1	2.0	-1.1	-11.5	5.3	1.8	-0.7
Non (cr)	1.9	-1.3	-0.2	0.3	-3.8	2.7	0.5	-0.8	-2.8	2.7	0.2	-0.5
Entry	n.a.				0.5				0.8			
					Medium Skilled Employment							
BM change:	-1.6				-7.0				-11.0			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	1.5	1.0	0.0	0.3	1.6	1.1	0.0	0.9	1.1	0.7	0.0	0.8
Med	-1.4	-3.7	-1.3	-2.5	-1.0	-3.8	-1.0	-3.0	-0.4	-2.4	-0.3	-2.9
High	0.0	1.7	3.0	1.7	0.1	2.7	3.9	2.1	0.1	2.4	2.7	1.3
Non	-7.1	10.6	-9.4	0.1	-10.2	16.8	-14.0	0.4	-8.0	18.2	-10.4	0.0
Non (cr)	-1.6	3.9	-1.8	0.0	-2.8	8.9	-3.3	0.3	-2.0	9.2	-0.9	0.0
Entry	n.a.				2.2				3.7			
					High Skilled Employment							
BM change:	1.8				6.2				9.2			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-0.3	0.0	-0.1	-0.2	-0.5	0.0	-0.2	-0.3	-0.4	0.0	-0.2	-0.3
Med	0.0	-2.4	-1.3	-1.6	0.0	-2.6	-1.1	-2.1	0.0	-1.6	-0.3	-2.0
High	0.3	1.7	3.4	1.9	0.5	2.7	4.5	2.4	0.5	2.4	3.2	1.6
Non	1.6	7.0	-10.5	0.7	3.5	11.8	16.0	1.6	3.5	12.9	-12.2	0.6
Non (cr)	0.4	2.5	-2.0	0.4	1.0	6.2	-3.7	1.1	0.9	6.6	-1.1	0.4
Entry	n.a.				2.1				2.9			

Contributions to employment share changes for older male workers based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A, B \in$ (Low, Med, High, Non, Entry). Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Older male workers comprise workers aged 51-65. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table 2.7: Counterfactual Results for Young Female Workers

Period	1981-1990				1981-2008				1981-2015			
	Low Skilled Employment											
BM change:	0.4				6.9				18.7			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	2.8	0.7	0.0	1.8	0.8	0.8	-0.1	0.7	1.2	0.8	-0.1	0.4
Med	-1.2	-3.7	0.0	-3.0	-0.5	-0.8	0.0	-2.5	0.7	1.1	0.0	-1.7
High	0.0	0.0	-0.3	-0.5	-0.1	0.0	-0.2	-0.7	-0.1	0.0	0.0	-0.8
Non	-3.5	4.0	0.9	0.0	-2.0	4.2	0.3	0.2	2.4	3.3	-0.4	-0.3
Non (cr)	0.5	0.1	0.0	0.0	1.7	-0.5	-0.7	-0.7	5.3	-1.6	-1.4	-1.4
Entry	1.4	0.5	0.4	0.0	6.8	4.7	0.5	-0.1	11.7	8.6	1.1	0.0
	Medium Skilled Employment											
BM change:	0.9				-8.4				-19.0			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-2.4	-0.7	0.0	-1.5	0.6	0.8	0.0	0.5	0.9	0.8	0.0	0.3
Med	1.2	5.1	0.0	3.9	-0.5	-1.1	0.5	-3.4	0.7	1.5	0.4	-2.1
High	0.0	0.0	-0.6	-0.9	0.0	0.2	0.3	1.0	0.0	0.4	0.0	0.7
Non	3.0	-5.3	1.7	0.1	-1.5	5.6	-0.5	0.1	1.8	4.1	0.4	-0.2
Non (cr)	-0.4	-0.1	0.1	0.0	1.3	-0.6	1.0	-0.4	4.0	-2.0	1.3	-0.9
Entry	-1.2	-0.6	0.7	0.0	5.5	5.7	-0.5	0.0	9.4	9.6	-0.7	0.0
	High Skilled Employment											
BM change:	-1.3				1.5				0.3			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	0.4	0.0	0.0	0.3	-0.2	0.0	0.2	-0.1	-0.2	0.0	0.1	-0.1
Med	0.0	1.4	0.0	0.9	0.0	-0.3	0.5	-0.9	0.0	0.3	0.5	-0.4
High	0.0	0.0	-0.9	-1.3	0.1	0.2	0.5	1.7	0.1	0.4	-0.1	1.5
Non	-0.5	-1.3	2.6	0.1	0.4	1.5	-0.8	-0.1	-0.6	0.8	0.8	0.1
Non (cr)	0.1	0.0	0.1	0.0	-0.4	-0.2	1.8	0.3	-1.3	-0.4	2.7	0.5
Entry	0.2	-0.1	1.0	0.0	-1.3	1.1	-1.0	0.0	-2.3	1.0	-1.8	0.0

Contributions to employment share changes for young female workers based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A, B \in$ (Low, Med, High, Non, Entry). Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Young female workers comprise workers aged 18-30. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

2.5 Discussion

I started the analysis referring to the notion that the process of creative destruction underlying job polarization does not need to make displaced workers worse off, as it also provides opportunities for job creation. As economic change may impact on workers not by diminishing the total number, but the type of jobs available, I argued one cannot infer this impact from distributional changes alone. Results from the counterfactual analysis allow relating changes in labour reallocation to job polarization, and so shed light on the impact job polarization had on workers.

Which changes in labour reallocation are associated with job polarization? Results call into question the narrative whereby medium skilled jobs disappear because redundant jobs are destroyed, and displaced workers reallocate to other jobs, low or high skilled. First, the decline in medium skilled employment is not driven by workers moving from medium skilled jobs to non-employment. Second, shares for low and high skilled jobs are not affected much by workers reallocating directly from medium skilled jobs or out of non-employment. Instead, medium skilled employment declines as workers who become non-employed, mostly from medium skilled jobs, do not return to medium skilled employment. Additionally, young workers directly sort into low skilled employment as they first enter the labour market.

What do these changes in labour reallocation imply about the impact at the worker level? That is, how does job polarization affect workers? For young workers, one can readily conclude that workers entering in low skilled jobs instead of medium skilled ones are worse off. It is, however, not clear whether workers are permanently worse off. Note that entering the labour market in low skilled jobs is not accompanied by increasing transition rates from low skilled jobs or non-employment to medium or high skilled jobs. This is suggestive that young workers are worse off not only by starting their career in a low skilled job, but also by not moving to better employment in later years. On the other hand, however, shifts towards low skilled employment are much stronger for young than for older workers across all periods. Young workers who age into older groups do not exhibit an increased propensity to hold a low skilled job. To fully understand employment trajectories for young workers after they enter in a low skilled job, one needs to go beyond analyzing annual transition rates. I leave this to future research. In any case, this is a novel finding in the literature. While US studies suggest that young workers have been disproportionately affected, they do not associate this finding with labour market entry rates.[48]

[48] See Cortes et al. [2014, 2016]; Smith [2013].

How about workers who remain non-employed, instead of returning to a medium skilled job? The fact that workers remain longer in non-employment allows for three separate explanations: first, workers seeking medium skilled employment spend more time searching because less medium skilled jobs are created. At the very least this would imply that workers are temporarily worse off, as they experience longer unemployment spells. Whether they are permanently worse off depends on what type of job they eventually return to. In this context, note that workers are still most likely to return to a medium skilled job, and for most of the sample period there is no general tendency for workers to be more likely to return to low or high skilled jobs. This may have changed more recently: in the most recent period young workers of both genders, previously employed in low or medium skilled jobs, more often leave for low and high skilled jobs, and prime aged workers of either gender leave more often for high skilled jobs. However, as these changes occur at the end of the sample period they have little effect on the counterfactual analysis.

Second, due to the lack of suitable employment prospects, non-employed workers may become discouraged and leave the labour force altogether. An increase in workers no longer seeking employment could reflect either changes in labour demand, e.g. jobs are available but under worse conditions, or changes in labour supply, e.g. changes in workers' outside options or preferences. As explanations for job polarization stress decreasing labour demand for medium skilled jobs, results would be suggestive for workers leaving the labour force as they no longer find suitable employment.

Third, workers may sort into types of employment which are mismeasured as non-employment in the dataset. This would apply to foreign workers who return home, or native workers who emigrate, or workers who become self-employed. As job polarization is an international phenomenon, it seems unlikely that medium skilled workers leave the UK to find better employment prospects elsewhere. Of course, emigration could have increased independently of job polarization. More plausibly, workers may have resorted to self-employment. Whether workers can be considered to be worse off under (part-time) self-employment depends on how they value self-employment over their alternative, medium skilled job.

NESPD and ASHE do not allow distinguishing non-employment from unemployment, inactivity, self-employment, or disappearance from the sample for other reasons, but it is possible to relate the results of the counterfactual analysis to population shares for employed, unemployed, and inactive workers for the same period and demographic groups, derived from the UK LFS. The corresponding graphs for male and female workers are shown in figures B.9 and B.10 in the appendix.[49] The counterfactual analysis implies that job polarization occurs in large parts as workers are non-employed rather than holding a medium skilled job, driven by a decrease in the outflow rate to medium skilled jobs. Also, remember that inflow rates do not contribute to, and if anything, counteract the decline. If workers are non-employed for longer periods because they take more time to search for jobs, one would expect the unemployment population share to increase. If instead workers move to inactivity, one should expect to see the inactivity population share rise. Finally, if workers move to self-employment instead of returning to medium skilled jobs (or they die or emigrate), there should be no fall in the employment population ratio.

Figure B.9 shows for male workers that LFS data are, by and large, compatible with medium skilled jobs disappearing because workers are more likely to be genuinely non-employed. Employment population shares decline for all age groups over the sample period, exhibiting particularly steep drops from 1980 onward. Employment population share for prime aged and older male workers reach their low point in the mid 90s, and recover slightly thereafter. In all cases, however, employment remains substantially below the level in 1975. Compared to 1975, young workers exhibit a decline of 10 pp, prime aged workers of about 5 pp, and older workers of as much as 20 pp. The long-term decline in employment population shares reflect a sustained increase in inactivity, while unemployment fluctuates over the sample period but does not rise permanently. This is compatible with workers being more likely to be inactive, and for some periods also unemployed, because they are less likely to return to medium skilled jobs. The implication is that job polarization may reflect workers not reallocating to other employment types but becoming inactive. One important question, of course, is whether these workers leave the labour force permanently, or only for short periods. Figure B.9 also suggests that the magnitude of shifts towards inactivity is largest for older workers, moderate for young and smallest for prime aged workers. This, of course, does not align with the relative importance of the decline in the outflow rate to non-employment, suggesting large shifts for prime aged workers relative to younger workers. The possibility remains that some of the contribution of the outflow rate, especially for prime aged male workers, reflects shifts to self-employment, or other types of mismeasurement, rather than genuine non-employment.

The non-employment outflow rate to medium skilled jobs is less important for women. While it counteracts the decline in medium skilled jobs for young women, it exhibits comparatively small contributions for prime aged and older women, which

[49]LFS data have been provided by Jennifer C. Smith (Elsby et al. [2011]).

are partly offset by declining non-employment inflow rates from medium skilled jobs. Figure B.10 is roughly in line with these findings, in that women do not exhibit a decline in inactivity similar to men. The employment share for all age groups increases, most strongly so for young female workers. This rise reflects a sustained fall in inactivity, beginning in the early to mid 80s. Unemployment, and in some cases also inactivity, rise initially, therefore leading to a temporary decline in the employment population share. The figure does not suggest shifts towards inactivity or unemployment. If such shifts occurred for prime aged and older women, they may have been compensated for by a general rise in the participation rate.

These results largely fit in with findings from the US. On the one hand the timing for UK job polarization differs substantially from the US. While the UK experienced job polarization since the mid 70s, the decline in routine manual employment started in the US in the mid 80s, and routine cognitive employment only declined since the 90s. On the other hand, studies using US data consistently find that the most important changes in transition rates accounting for job polarization are declining rates from non-employment to medium skilled or routine jobs. These studies suggest that outflow rates from both unemployment and non-participation are important.[50] Moreover, Cortes et al. [2016] find that job polarization largely reflects workers shifting from routine employment to non-employment. Interestingly, they find that low educated male workers are more important for the decline in routine manual jobs, while women with medium levels of education are more important for the decline in routine cognitive employment. Thus, while the timing, and possibly the demographic groups affected most, may differ, findings for both the US and UK suggest that job polarization reflects a decline in demand for medium skilled jobs, reallocating workers from medium skilled jobs to non-employment via a fall in the job finding rate for medium skilled jobs.

To summarize, the analysis suggests that job polarization does not primarily impact on workers by reallocating them directly from medium to either low or high skilled jobs. Instead, to a large extent the disappearance of medium skilled jobs reflects workers spending more time non-employed. This may reflect workers being unemployed for longer, or workers may permanently leave the labour force, or workers may have migrated or sorted into self-employment. Results from the US are suggestive that transitions into unemployment and non-participation are most important.[51] Assuming this is the case, an important question for future research is why changes in firms' demand for job types result in workers being

[50]See Cortes et al. [2014, 2016]; Foote and Ryan [2014].

[51]See Cortes et al. [2016] for an exposition of a model explaining a worker's decision to move to non-routine manual jobs or non-participation in presence of automation.

unemployed or leaving the labour force, rather than in workers adapting to the change in demand. Barriers in educational attainment may play an important role, as findings suggest that less educated worker groups are generally most adversely affected, and that more educated workers are more likely to move to high skilled jobs.[52] Additionally, many young workers are worse off as they reallocate to low skilled jobs or start their career in lower quality jobs. Overall, these findings hold the potential that job polarization may have had a large and adverse impact on many workers. To fully understand the impact of job polarization, one needs to go beyond the analysis of annual transition rates, to understand whether young workers starting their careers in low skilled jobs eventually upgraded to better jobs, and whether longer non-employment spells, implied by the decline in non-employment outflow rates to medium skilled jobs, reflects workers remaining unemployed for longer, taking more time to reallocate to other jobs, or whether they become discouraged and leave the labour force. To better understand the link between education and the impact of job polarization, one needs to examine the importance of transition rate changes on job polarization conditional on education.[53]

[52]See Cortes et al. [2016] for showing that low and medium educated workers are most affected by job polarization in the US, Salvatori [2015] for showing that non-graduates can account for the entire decline in medium skilled jobs, and that graduates did not experience job polarization in the UK. See Foote and Ryan [2014] and Upward and Wright [2007] for evidence that more high skilled workers are more likely to transition from medium to high skilled jobs.

[53]Unfortunately, while this was not possible with the dataset used for this analysis, the British Household Panel Survey and the Quarterly Labour Force Survey both contain information on job polarization and transition rates, while allowing controlling for education. These disadvantage of these datasets is that they only start in the 1990s and so cover substantially shorter periods.

Chapter 3

Gone for Now or Gone for Good – On the Impact of UK Job Polarization on Non-Employment Duration

3.1 Introduction

Over at least the last four decades the number of medium skilled jobs declined in the UK. That is, jobs paying moderate wages disappeared, in absolute terms and relative to jobs paying either very low or very high wages. This process of declining medium skilled jobs, or jobs intensive in routine tasks, and the simultaneous increase in low and high skilled jobs, or jobs intensive in non-routine non-cognitive and non-routine cognitive tasks, is referred to as job polarization. Medium skilled jobs disappeared because, first, globalization allowed these jobs to be performed more cheaply by workers abroad, and, second, advances in information and communication technology (ICT) capital made it possible to employ capital substituting for tasks previously done by medium skilled workers.[1]

One concern raised in this context is how the decline in medium skilled employment impacts on workers who would otherwise be employed in such jobs. The notion that technological change, or globalization, makes workers worse off by destroying their jobs, referred to as technological unemployment, is countered by the argument that the resulting higher labour productivity provides incentives to create

[1]For a summary of the literature on job polarization see Acemoglu and Autor [2011]. For a more recent overview see Autor [2014] and Autor et al. [2015].

new jobs. As redundant medium skilled jobs disappear, workers reallocate to new types of jobs.[2] The literature examining worker flows underlying job polarization in the US, however, argues that the disappearance of medium skilled jobs to large parts reflects shifts in the distribution of labour market states towards non-employment. Chapter 2 shows that similar results obtain for the UK. This shift towards non-employment seems to be driven by the outflow rate to medium skilled jobs: to a large extent fewer workers are employed in medium skilled jobs because they are non-employed for longer periods.[3]

Does the increase in non-employment duration imply that workers take longer time to reallocate, or does it reflect workers failing to reallocate and become permanently jobless? An increase in non-employment duration associated with job polarization is compatible with two scenarios: In the first case, job polarization leads to a decline in medium skilled jobs, but does not render workers permanently jobless, as eventually they reallocate to new jobs. Workers suffer longer periods without jobs because they take longer time to reallocate. Non-employment spells are longer, but still temporary. Longer non-employment spells in this case may reflect, for instance, workers receiving fewer job offers or searching less intensively. In the second case, job polarization leads to a decline in medium skilled jobs and, as a result of the worsening of employment prospects, workers do not reallocate to new jobs and so become permanently jobless. Assuming that workers are eventually able to find jobs at the going wage rate, the failure to reallocate suggests that the decline in medium skilled jobs leaves workers with job opportunities which are either too unattractive or unavailable to them. Workers stop searching for employment altogether. Unless they start searching again, they fail to reallocate to new jobs and become permanently non-employed. After all, job polarization may result in technological non-employment as it replaces ordinary, medium skilled jobs with either low skilled 'lousy' jobs, or with high skilled, but for many unattainable, 'lovely' jobs. Given these alternatives, workers unable to move up the job ladder, and unwilling to move down, may end up being permanently jobless.

Previous work established that job polarization reflects workers becoming more likely to be non-employed and less likely to hold a medium skilled job because of a decline in the outflow rate to medium skilled jobs. The datasets used for these papers generally follow workers for up to 1 year, however, and are therefore unable to address heterogeneous implications on non-employment durations.[4] Do longer, average non-employment durations reflect all workers experiencing longer

[2] See, for instance, Bartelsman [2013] and Autor [2014].
[3] See Cortes et al. [2014, 2016]; Smith [2013]
[4] See Cortes et al. [2014, 2016]; Smith [2013].

non-employment spells, as in the first scenario, or do they reflect some workers becoming jobless for very long periods, as in the second scenario? In this chapter I address the question whether job polarization is associated with an increase in longer but temporary, or permanent non-employment spells, which in turn is suggestive for job polarization leading to workers taking more time to reallocate, or failing to reallocate altogether.

Addressing this question requires distinguishing short-term, temporary, and permanent non-employment spells. Adopting a worker flow perspective, workers may return to employment eventually, and yet be said to have failed to reallocate. A permanent spell is very long, suggesting hardship on the worker's side, but not infinite. To classify non-employment spells as short-term, temporary, or permanent, the main analysis adopts the following categorization: Short-term spells last less than one year, temporary spells last at least one year but less than five years, and permanent spells last five years or more.[5] I take it that classifying spells lasting at least five years as permanent is a conservative measure. It seems implausible that workers being non-employed for such long periods are actively searching for jobs throughout this period, or expect to return to employment anytime soon. It is in that sense that I refer to workers who do not return to employment for at least five years as having failed to reallocate.

Additionally, the question requires relating changes in non-employment duration to job polarization. For the UK, job polarization can be linked to an increase in average non-employment duration, as the decline in the job finding rate to medium skilled jobs can account for a substantial part of the defining decline in medium skilled jobs: everything else constant, a decrease in the job finding rate to medium skilled jobs leads to longer average non-employment duration. Exit rates to other jobs may change, however. Associating changes in observed non-employment with job polarization may be misleading. I therefore employ survival analysis to identify changes in non-employment duration due to changes in the job finding rate to medium skilled jobs, which in turn are associated with job polarization.

In particular, using a flow sample approach based on UK panel data covering the period from 1975 to 2015, I estimate a discrete-time non-parametric competing risks model, allowing for exits from non-employment to low, medium, and high skilled jobs, and decompose changes in the distribution of non-employment spell durations into changes in job-specific hazard rates. This allows investigating to what extent a change in the fraction of workers experiencing temporary or perma-

[5] I also report results separating temporary from permanent spells at shorter and longer durations as a robustness check.

nent spells reflects a change in the exit rate to medium skilled jobs, i.e. whether job polarization is associated with an increase in temporary or permanent non-employment spells. Additionally, the analysis reveals to what extent workers avoid longer non-employment spells by reallocating to other job types, i.e. whether exits to low or high skilled jobs exacerbate or dampen the effect of job polarization on non-employment duration.

Chapter 2 has shown that the extent of job polarization and the underlying worker flows differ substantially among demographic groups. Given this heterogeneity, and to relate findings to this literature, I conduct the analysis at the demographic group level, considering gender groups and three age bands. To focus the analysis, I consider only spells starting in expansionary periods, i.e. I compare non-employment spell durations across different periods for workers who left work during 'normal' times. Reasons for entering non-employment, and accordingly the time spent in non-employment, are likely to differ during expansions and recessions.[6] Spells of workers entering non-employment during a recession therefore warrant special attention. I leave this analysis to future research.

Understanding whether job polarization implies longer reallocation periods, or whether it renders some workers permanently jobless, is important because it contributes to our understanding of how costly job polarization has been for workers, and it is telling about the nature of these costs. Understanding the magnitude and source of these costs is of direct importance for policy makers. For instance, results are suggestive for whether job polarization led to an increase in unemployment duration, as workers took longer time to reallocate, or to an increase in inactivity, as workers ceased searching and became permanently jobless. A policy implication in case of increased unemployment duration may be to facilitate labour reallocation by encouraging job creation at the firm side, or search effort or retraining at the worker side. Results may be informative about the applicability of potential policies. Conducting the analysis at the demographic group level additionally allows identifying affected worker groups requiring particular attention. Results may also be suggestive whether changes in non-employment duration apply to all groups equally, and therefore likely reflect aggregate factors independent of individual worker be-

[6]Labour market flows are generally recognized to vary systematically between recessions and expansions. A sizeable literature focuses on cyclical variation in unemployment inflow and outflow rates. See, for instance, Elsby et al. [2011]; Gomes [2012]; Elsby et al. [2015]; Shimer [2012]; Fujita and Ramey [2007]. Haltiwanger [2011] argues for the US that earnings losses associated with displacement are more severe during recessions, and non-employment duration increases, consistent with the view that worker separations during expansions are more likely to be associated with workers reallocating from low-productivity to high-productivity firms. Cyclical fluctuations in job-to-job transitions and recall have been examined extensively in the US. See Fallick et al. [2012]; Moscarini and Postel-Vinay [2012]; Fujita and Moscarini [2013].

haviour, or whether they differ across groups, so there is scope to avoid longer spells by affecting workers' incentives. To the extent job polarization results in permanent joblessness, it is important to understand why workers did not reallocate. For instance, does reallocation take longer because workers are unaware of job opportunities in alternative job types? Do they choose not to reallocate because of their unwillingness to accept lower paying or lower skilled jobs? Could those workers rely on other sources of income, such as other employed household members? To what extent did workers fail to reallocate because they were unable to gain qualifications for other job types?

For male workers, I find evidence that job polarization led to more workers taking longer time to reallocate, and that some workers became permanently jobless due to the decline in demand for medium skilled jobs. For female workers, results are more mixed and not generally suggestive for job polarization having adversely affected womens' prospects to exit non-employment quickly. These differences across demographic groups are suggestive that the interaction of worsening aggregate labour market conditions and individual behaviour, correlated with demographic characteristics, led to the observed decline in non-employment exit rates. The question arises what these differences in behaviour were. In any case, they suggest that a decline in demand for medium skilled jobs does not necessarily have to result in longer non-employment spells.

The remainder of this chapter proceeds as follows. I discuss the construction of the data in section 3.2. Section 3.3 presents non-parametric estimates of the survival function and discusses how the change in aggregate transition rates from non-employment is linked to changes in the distribution of survival times. Section 3.4 introduces the competing risks model, and presents results for the decomposition of the overall survival function. Section 3.5 discusses.

3.2 Data

As in chapter 2, I combine NESPD and ASHE to cover the years 1975 to 2015, restricting the sample to observations for workers aged 18 to 65 years.[7] I use the same procedure described in chapter 2 to create a consistent variable categorizing workers' occupations according to SOC 2010 codes over the entire sample period. Based on the distribution of occupational median wages in 1975, I create categories for low, medium and high skilled occupations, again following the procedure described in chapter 2. I then group each worker/year observation into one of four mutually

[7]See ONS [2017d,a].

exclusive labour market states, i.e. workers are categorized to be employed in a low, medium or high skilled job, or to be non-employed. Having done so, I restrict the sample to non-employment spells of workers.[8]

Information about non-employment is inferred in the sample as follows. The sampling procedure of ASHE and NESPD implies that workers once included in the sample will occur in every year in which they have been registered with PAYE during the reference period in April. Workers disappear from the sample because of non-employment, self-employment, death, or emigration. I treat workers as non-employed whenever they are not recorded as employed and are aged between 18 and 65 years. This suggests that observations for non-employment are suffering from measurement error, i.e. some workers who are observed as non-employed in fact left the country, or entered self-employment, or died. Unfortunately, it is not possible to avoid this measurement error using this data. However, as I compare changes in non-employment duration over time, results will only be affected to the extent the frequency with which workers died, emigrated, or entered self-employment changed over time.

I distinguish six demographic groups: three age groups and gender groups. Age groups are defined for age bands of 18-30, 31-50, and 51-65 years. For each demographic group I create a sub-sample containing all non-employment spells for which the worker belongs to the respective demographic group at the time he or she is first recorded as non-employed.

Additionally, I split the overall sample period into seven either expansionary or contractionary periods, following the classification described in chapter 2. Periods 1, 3, 5, and 7 are expansionary periods, referring respectively to years 1975-1978, 1982-1989, 1993-2007, and 2010-2015. Periods 2, 4, and 6 are contractionary periods, referring to years 1979-1981, 1990-1992, and 2008-2009.[9]

To organize the data for the survival analysis, I first create identifier variables for each spell. For each spell i in the respective sub-sample, I then construct the following variables: I create a set of dummy variables denoted δ_i^p for $p = 1, 2, \ldots, 7$, indicating the period in which the non-employment spell starts. I create variable d_i for the spell duration, measured in completed years of observed non-employment, and ranging from zero in the first year the spell is observed until K_i, the last year the spell is observed. One can think of the resulting dataset as panel, whereby the

[8]A spell contains all consecutive observations whenever a worker is recorded as non-employed and, if the worker returns to employment, the first observation following the end of the current non-employment spell.

[9]This accords with the definition of expansionary and recessionary periods used for transition rates in chapter 2: Expansionary periods start the year after the through of the preceding recession and end the year before the peak of the subsequent recession.

spell identifier identifies the group, and duration replaces the time variable. Overall, this gives 183,204 spells for young women, 221,443 spells for prime aged women, and 117,859 spells for older women. For young men, it gives 179,653 spells overall, for prime aged men it gives a total of 241,855 spells and for older male workers it gives a total of 150,133 spells.[10]

Each sub-sample is a flow sample of spells starting in the respective period and demographic group. The advantage of flow sampling is to avoid the bias associated with stock sampling. Stock sampling, e.g. including all spells which are active during a given period, may bias exit rates downwards, and thus survival functions upwards, as workers with low exit probability will comprise a larger share compared to a flow sample. In the particular case of the dataset at hand, flow sampling additionally allows avoiding the sample start problem described in chapter 2. Finally, I create a set of dummy variables indicating at each duration whether the spell ends with return to low, medium, high skilled employment or overall employment.[11] Spells are censored if the worker is still non-employed at the last observation of the spell. In the sample spells are censored if workers are non-employed at age 65, as they exit the sample from the following year onwards, or if workers are non-employed in 2015. Lastly, note that spells are treated as independent.[12]

I refer to spell duration as the completed number of years of observed non-employment with reference to the following limitations.[13] First, I refer to observed years of non-employment as I only observe the worker's employment status at one

[10]The number of spells for each period varies with period length. However, each period contains a large number of spells. For young women, the number of spells for each period ranges from 19,079 in period 1 to 86,330 in period 5. For prime aged women, it ranges from 19,013 in period 1 to 114,395 in period 5, and for older women from 12,415 spells in period 1 to 54,823 in period 5. For young men it ranges from 22,279 spells in period 1 to 80,342 in period 5, for prime aged men from 27,018 spells period 1 to 118,161 in period 5, and for older male workers from 22,050 spells in period 1 to 63,419 in period 5.

[11]The implementation of survival analysis requires data are organized in the correct form. When computing Kaplan-Meier estimates, I keep only the last observation for each spell.

[12]I do not account for repeated spells of the same worker. Doing so would allow accounting for unobserved heterogeneity. However, controlling for unobserved heterogeneity substantially complicates estimation in a competing risk framework. Unless strong assumptions are imposed, e.g. independence of unobserved heterogeneity across destinations, standard statistical programmes cannot be used. Omitting unobserved heterogeneity can bias hazard rates downwards at longer durations. As my interest lies in identifying changes in survival functions over time, rather than in identifying the correct shape of the hazard function, I ignore unobserved heterogeneity.

[13]Note the following remarks for the interpretation of spell duration: ASHE and NESPD contain annual, point-in-time observations for a reference period in April of the respective year. A non-employment spell starts in the first year a worker is observed as non-employed with duration zero. If the worker is still recorded as non-employed in the following year, one completed year of non-employment is observed. Accordingly, if the worker is still non-employed in the xth year after the non-employment spell started, I observe x completed years of non-employment, and the duration of the spell at this point in time is x. If a spell ends at duration x with the worker being observed as employed, I observe $x - 1$ completed years of non-employment.

point in time in April each year. I do not observe whether the worker returned and re-entered non-employment in between two subsequent years in which the workers is recorded as non-employed. Second, I refer to completed years of non-employment as, for the same reason, I do not observe the exact entry and exit date of the non-employment spell. Suppose for instance a worker enters non-employment in May 2005 and exits in May 2006. This worker would thus be observed as non-employed in 2006 and as employed in 2007, implying the spell duration is zero: zero completed years of non-employment are observed, although the actual but unobserved spell duration exceeds 1 year.

Finally, a further limitation arises as I only observe non-employment spells at annual frequencies. I miss short-term non-employment starting and ending in between survey dates, e.g. workers entering non-employment in May 2005 and exiting in June 2005. This time aggregation is known to introduce a downward bias for the level of non-employment outflow rates.[14] As I focus only on expansionary periods, effectively comparing changes in the level of outflow rates across periods, I argue that the time aggregation bias is not of major concern for my analysis.

These limitations are likely to affect results as follows. Note that one may think of the data limitations as inducing measurement error for the level of spells with different durations. For instance, spells which are observed with duration 1 actually reflect two short spells with duration zero, or spells with duration zero actually reflect longer spells. Also, some non-employment spells do not reflect true non-employment, but workers becoming self-employed, emigrating or having died. Measurement error may be systematic. The fact that long spells can reflect several short spells suggests that the sample overstates the level of spells with longer durations, although this may itself become less severe with longer observed spell duration. Additionally, as workers with longer durations are more likely to occur in the point-in-time sample it also likely understates the level of short duration spells. Some of the spells with duration zero reflect spells which begin some time before April, when they are first observed, and last for at least one year. However, as I focus on changes in the frequency with which spells of different durations occur, to the extent systematic measurement error only occurs in levels it will be differenced out, and observed changes in the incidence of spells with different durations will not be affected systematically. These data limitations affect results for changes in the incidence of spells only to the extent that measurement error varies systematically with changes in spell durations. Arguably, this will be most severe with regard to

[14] See Shimer [2012] for a discussion of time aggregation. As mentioned before, time aggregation is of concern for the literature on cyclical unemployment flows because the bias is likely to vary with the cycle.

workers moving to self-employment instead of non-employment. It is plausible that self-employment increased over time as a result of the worsening of labour market prospects for workers seeking medium skilled jobs. This would be systematically related with increases in the incidence of spells with long durations, assuming that workers spend some time in self-employment. Also, it is plausible that globalization led to an increase in workers migrating to and out of the UK, again suggesting that the incidence of long-term non-employment spells may be overstated.[15]

3.3 Changes in Distribution of Non-Employment Duration

This chapter examines whether job polarization is associated with workers being more likely to experience longer temporary or permanent non-employment spells. In chapter 2, I established that UK job polarization largely reflects a shift from medium skilled jobs to non-employment, brought about by workers being less likely to leave non-employment to medium skilled jobs. This implies that the decline in the non-employment outflow rate to medium skilled employment led to longer non-employment spells. Longer non-employment spells, in turn, can reflect either longer temporary or more permanent spells.

In this section, I first demonstrate how the decline in the exit rate to medium skilled employment led to an increase in average non-employment duration. This motivates the analysis examining how changes in the exit rate to medium skilled employment affect the distribution of non-employment spell durations. The distribution of spell durations is depicted in terms of the survival function, stating the fraction of workers remaining in non-employment at each duration.[16]

Second, I present non-parametric estimates of survival functions for each period and group. These survival functions provide a largely unadulterated account of the evolution of the distribution of non-employment durations over the sample period. I present these to discuss to what extent lower annual exit rates are reflected in longer durations for spells starting in later periods. This serves to show that the

[15] I discuss below changes in group-specific population shares for employment, inactivity and unemployment, which shed some light on the plausibility of results reflecting genuine non-employment.

[16] Note the following comments on terminology: I refer to transition rates unconditional on non-employment duration as annual transition rates or non-employment outflow rates. This is the number of workers in a given year who move to some type of employment in the following year, relative to all non-employed workers in that year. I refer to duration-specific exit rates as exit or hazard rate. This is the number of workers who are non-employed at a given duration who exit to some type of employment relative to the number of workers non-employed at that duration. See also the formal discussion of hazard rates below.

decrease in aggregate annual transition rates in subsequent periods implies longer non-employment spells for spells starting in later periods.[17] I also use these figures to discuss more broadly how the distribution of non-employment spell durations changed. This shows how the incidence of short-term versus temporary or permanent non-employment spells changed over the sample period. The subsequent analysis examines how these changes are related to job polarization.

3.3.1 Changes in Average Non-Employment Duration

Figure 3.1 shows annual transition rates from non-employment to medium skilled employment for male and female workers by age group. The figure shows both the uncorrected annual transition rates as well as corrected period transition rates for the first two periods.[18]

Overall, the figures demonstrate a widespread and prolonged decline in the probability to leave non-employment to medium skilled jobs for all demographic groups except for young women.[19] The decline in exit rates generally occurs in the early years of the sample period, the late 70s and early 80s, but in some cases continues well into the 90s. Following the decline, outflow rates remain largely flat until the end of the sample in 2015. Male workers generally experience a larger decline in the probability to leave to medium skilled jobs than women. Qualitative conclusions remain largely unchanged when using corrected transition rates instead of uncorrected ones. Corrected rates generally imply a smaller decline in exit rates. This is to be expected as the sample start problem would imply artificially large outflow rates in early years. However, even with corrected transition rates, we observe a drop in these rates from the first and second period onward. Young women represent the exception. Using corrected rates, outflow rates first remain fairly constant and increase to an all time high in the most recent years. The previous

[17] The two concepts are related as longer survival times imply lower transition rates per se. However, they are dissimilar as non-monotonic changes in survival functions may not necessarily imply a decrease in transition rates, and to the extent that the lengths of spells starting in one period partly reflect annual transition rates in later periods.

[18] As mentioned in chapter 2, non-employment outflow rates may decline right after the sample start, because of either worker heterogeneity or duration dependence leading to changes in the composition of non-employed workers. This composition reflects that workers at different durations exhibit different outflow probabilities. I address this sample start problem by correcting non-employment outflow rates in the first two periods, from 1976 to 1981, using the composition of workers by duration spent in non-employment observed for period 3, from 1982 to 1989. Any decrease in corrected rates from period 1 and 2 to period 3 only reflects changes in duration-specific exit rates, which are independent of the sample start problem. While the uncorrected transition rates can be seen as an upper bound for actual transition rates, i.e. in case the initial conditions problem does not apply, corrected transition rates constitute a conservative measure and can be seen as a lower bound for transition rates from non-employment in these periods.

[19] Note that the jumps in 2007 and 2008 are unrelated to any known discontinuities.

chapter established that the decline in these rates can account for a substantial part of job polarization. In particular, depending on using corrected or uncorrected rates, this decline can account for half up to the entire decline in the medium skilled employment share.

The implications of these changes are that workers of all demographic groups, except young women, experience longer non-employment spells on average because they are less likely to leave non-employment to a medium skilled job. Whether actual non-employment duration increases overall, however, depends not just on the probability to leave to medium skilled jobs, but on the probability to leave to any type of employment. To show that the probability to employment of any type decreases as well, figure 3.2 shows the overall outflow rate from non-employment by demographic group. The overall outflow rate is the sum of the probability to leave to low, medium or high skilled employment. I again show corrected and uncorrected rates.

Changes in the overall exit rate appear to be largely driven by changes in the outflow rate to medium skilled employment. Both rates exhibit similar changes over time, i.e. a sharp drop from the mid 70s onwards, and a levelling off at a permanently lower level, latest in the early 90s. This implies that, as workers leave to medium skilled jobs less frequently, they do not generally become substantially more likely to leave to other job types. The decline in exit rates to medium skilled jobs is evidently not offset by changes in exit rates to low and high skilled jobs. Some notable differences remain. First, workers of all age groups see an increase in their exit probability in the most recent period, suggesting that workers generally become more likely to leave to other job types in later years. Second, young male workers exhibit a smaller decline, and young women a larger increase, in exit rates to overall employment than to medium skilled employment. This suggests that these workers become more likely to leave to other job types already in earlier years.

Overall, we observe that the decreasing probability to leave non-employment to a medium skilled job is accompanied by workers being less likely to leave to employment in general. This readily implies that, on average, workers experience longer non-employment durations: the expected duration in non-employment is just the inverse of the exit rate. To see this more clearly, figure 3.3 shows for each group the period average non-employment duration, computed as the inverse of the non-employment overall outflow rate averaged over each period. Panel (a) of figure 3.3 gives results using corrected rates, and panel (b) shows for uncorrected outflow rates.[20]

[20]Time aggregation implies that non-employment outflow rates may be understated. Due to

Both figures demonstrate a widespread, and in most cases lasting, increase in average non-employment duration. Average non-employment duration generally peaks in periods 4 or 5, lasting from 1990 to 1992 and 1993 to 2007 respectively, and while duration becomes shorter in subsequent periods, in most cases it remains above the level observed at the sample start. This pattern reflects exit rates mostly falling throughout the 80s until the 90s and the subsequent levelling off, driven by the decline in exit rates to medium skilled jobs. The relative shorter average non-employment duration in later periods arguably reflects workers being more likely to leave to other than medium skilled job types. Young workers stand out, if using corrected rates, as the smaller decline in exit rates to overall employment for young male workers suggests only a temporary increase in average non-employment duration, and young women exhibiting an absolute increase in exit rates to employment implies ever shorter average durations.

The magnitude is generally large in absolute terms, implying longer average durations of several years, and it is generally largest for prime aged and older male workers, and less pronounced for women. In particular, average duration for prime aged male workers increases from about two to three years (four years for corrected rates) in the first two periods to around seven years in period 5, settling at just below six years in the most recent period. For older workers, the increase is even larger, and ranges from about two to three years (four years for corrected rates) to close to eight years in period 5. For younger male workers duration still increases from about two (three for corrected rates) to four years in period 4, before returning close to its initial value. For prime aged and older women, results suggest temporary increases in average duration of around two years when considering corrected rates, peaking in period 5 and decreasing thereafter to levels close to those for corrected rates in periods 1 and 2. Using uncorrected rates, the rise is up to four years, and rates do not return to their initial level in period 7. As uncorrected rates are likely to overstate exit rates in periods 1 and 2, a cautious interpretation of results suggests evidence only supports a temporary increase in average duration for these women.

To sum up, there is conclusive evidence that average non-employment duration increases by a magnitude of several years for most groups, especially for prime aged and older male workers. Workers generally face the longest average non-employment duration during the 90s, and most groups experience permanently longer durations. For young male and prime aged and older female workers there is more mixed evidence, possibly suggesting only a temporary increase in average

annual observations I miss some short non-employment spells with high exit probability. The level of average non-employment duration may be overstated as a result. Conclusions in this section are largely unaffected as I compare changes in the level across periods.

non-employment duration. Young women experience at most a modest temporary increase, and arguably face shorter non-employment durations in recent periods compared to the first period.

3.3.2 Changes in Survival Functions

As seen above, non-employment exit rates suggest that most worker groups, on average, experience longer non-employment spells. These exit rates implicitly assume all workers face the same probability to leave non-employment, regardless of the duration already spent in non-employment. A worker non-employed for ten years faces the same decrease in the probability to leave non-employment as a worker just entering non-employment. On this account, a decrease in average non-employment duration suggests a rightward shift in the distribution of non-employment spell durations. That is, all workers experience an increase in non-employment duration by the same amount.

This ignores the possibility that workers experience differential changes in non-employment duration. Allowing for this possibility lies at the heart of this chapter, as the distinction between job polarization leading to workers taking more time to reallocate or failing to reallocate and being permanently jobless draws upon differential changes in the probability of experiencing spells of particular durations.

To allow for this possibility, I now turn to survival analysis, which explicitly accounts for duration dependence. The relationship between non-employment duration and the probability to leave to non-employment, expressed in terms of a hazard rate, is estimated based on spells, which possibly last over various periods. This necessitates a conceptual deviation from outflow rates: period average non-employment outflow rates reflect a lower probability to leave non-employment in a particular period, while hazard rates reflect a lower probability to leave non-employment at a given duration for spells which begin in a particular period.

The aim of this sub-section is to show how these two concepts are related for the issue at hand. I want to show that the increase in average non-employment duration – implied by the decrease in the non-employment exit rate – corresponds to longer spell durations for spells starting in later periods, relative to spells starting in period 1, in terms of changes in the overall survival function. Given this conceptual difference, it is important to demonstrate this correspondence in order to establish the plausibility of relating job polarization, which is linked to decreasing non-employment outflow rates to medium skilled jobs, to changes in survival functions. Additionally, relating changes in average non-employment duration to survival functions is informative about the implied changes in the distribution of

Figure 3.1: Non-employment outflow rate to medium skilled jobs

(a) Male workers

(b) Female workers

Annual transition rates from non-employment to medium skilled employment for male and female workers from 1976 to 2014 by age group. Suffix (cr) indicates transition rates averaged from 1976 to 1978, and 1979 to 1981, corrected for sample start problems using spell duration composition for years 1982 to 1989. Based on annual observations from NESPD and ASHE.

Figure 3.2: Non-employment outflow rate to employment

(a) Male workers

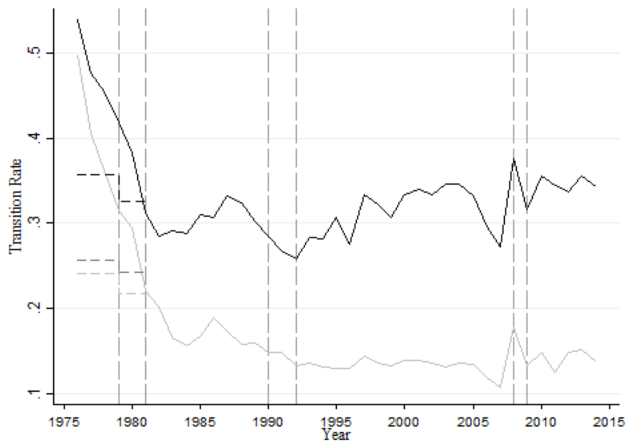

(b) Female workers

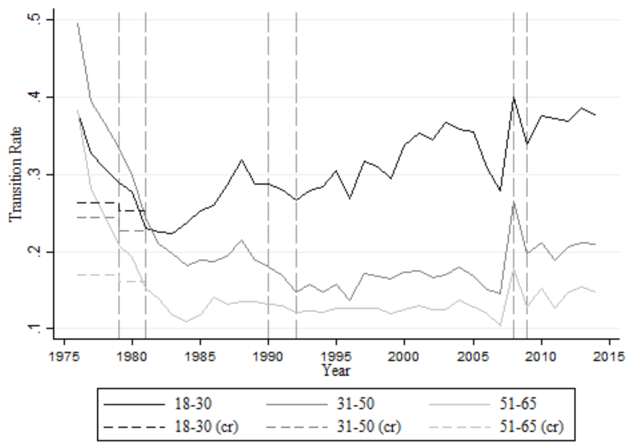

Annual transition rates from non-employment to employment for male and female workers from 1976 to 2014 by age group. Suffix (cr) indicates transition rates averaged from 1976 to 1978, and 1979 to 1981, corrected for sample start problems using spell duration composition for years 1982 to 1989. Based on annual observations from NESPD and ASHE.

Figure 3.3: Average Non-Employment Duration

(a) Corrected

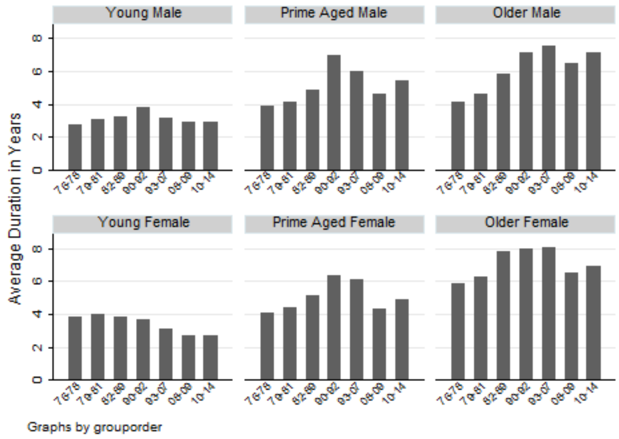

(b) Uncorrected

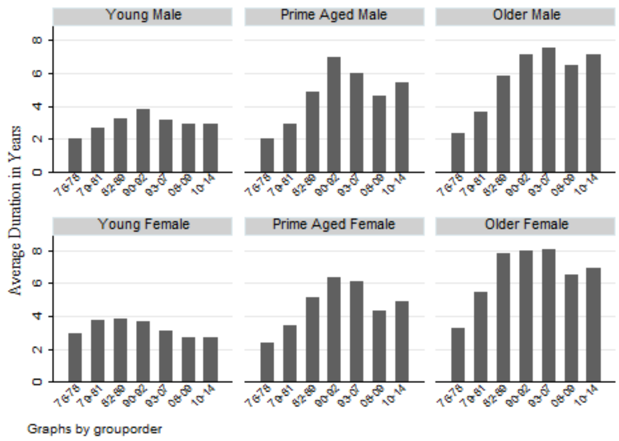

Average period non-employment duration for male and female workers from 1976 to 2014 by age group. Average non-employment duration for a particular period is the inverse of annual transition rates averaged over respective period. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Panel (a) based on corrected rates. Period transition rates from 1976 to 1978 and 1979 to 1981 have been corrected for sample start problems using spell duration composition for years 1982 to 1989. Panel (b) based on observed annual transition rates. Based on annual observations from NESPD and ASHE.

non-employment spell durations, i.e. whether workers experience generally longer spells or whether changes are concentrated among some workers becoming permanently non-employed. Lastly, providing evidence for longer spells based on survival functions provides additional evidence for a decrease in non-employment exit rates which is unaffected by the sample start problem.

To do so, I examine non-parametric estimates, i.e. Kaplan-Meier estimates, of the survival function for non-employment spells with regard to exits to overall employment. Kaplan-Meier estimates for the survival function are constructed as follows[21]: for each subsample of spells starting in period p for workers in demographic group g, denote the number of spells with exit recorded at duration j as h_j. Denote the number of spells censored at duration j as m_j. Recall that spells are censored if spells are still ongoing at the sample end or if workers are non-employed when reaching retirement age.[22] Further, denote the number of spells neither completed nor censored before duration j, i.e. the number of active spells or the population at risk at j, as n_j, i.e. $n_j = \sum_{i \geq j}^{K}(h_i + m_i)$, where K denotes the maximum observed duration of all spells in the subsample.[23] Denote the hazard rate at duration j as λ_i. This is the probability to leave to employment at duration j conditional on surviving until duration j. The Kaplan-Meier estimator for the hazard rate at duration j is then computed as:

$$\hat{\lambda}_j = \frac{h_i}{n_i} \tag{3.1}$$

For instance, consider $\hat{\lambda}_1 = h_1/n_1$. In this case h_1 is the number of spells for which the worker is recorded as non-employed at duration 0 and as employed at duration 1. The population at risk, n_1, is the total number of spells. The hazard rate at duration 1 is equal to the fraction of spells, for which the observed number of completed years of non-employment is 0, of the total number of spells. Denote

[21] See Kiefer [1988]; Allison [2014]; Jenkins [2005] for introductions to survival analysis, and discussions of non-parametric estimators for hazard and survival functions in particular.

[22] In principle, one could treat retirement as exit, so that the survival function gives the fraction of workers remaining non-employed at a particular duration net of workers who left to retirement. However, retirement is defined purely in technical terms and does not reflect any economic event. As exits to retirement do not vary systematically across periods, results for changes in survival functions across periods are not affected by including or excluding exits to retirement. In order to account for exits to retirement, for the overall survival function exits to retirement would be treated equivalently to workers returning to employment, i.e. hazard rates would be higher. My focus is on changes in survival functions due to changes in exit rates to employment. I therefore abstract from exits to retirement and treat these as censored.

[23] Differences in K imply differences in the domain for which survival functions can be computed. K varies by subsample because of differences in the age at which workers can enter non-employment relative to the fixed retirement age of 65, and because spells starting in later periods are closer to the sample end in 2015.

the survival function for duration j as S_j, corresponding to the probability of remaining in non-employment at duration j. For instance, the value of S_1 corresponds to the estimated probability of experencing a non-employment spell with at least duration 1; or equivalently the fraction of spells for which one completed year of non-employment is observed, i.e. which have been recorded as non-employed in two subsequent years. The estimator for the survival function is then computed as:

$$\hat{S}_j = \prod_{i=1}^{j}(1 - \hat{\lambda}_i) \tag{3.2}$$

As observations about remaining workers are made at discrete time intervals, the survival function takes the form of a step-function. For instance, $\hat{S}_1 = (1-\hat{\lambda}_1) = (1-h_1/n_1)$ is the fraction of spells still active at duration 1, i.e. the fraction of spells for which the observed completed number of non-employment spells is at least 1.

I want to show how survival functions changed over time. To do so, I compare survival functions for spells starting in different expansionary periods. I focus on spells that begin in expansionary periods as I want to focus on changes in non-employment duration for workers who left employed under 'normal' circumstances. As argued above, spells for workers who enter during a recession are likely to be different. Also, it has been argued for the US that job polarization would be related to jobless recoveries, whereby medium skilled workers would be laid off during recessions without returning to similar employment afterwards.[24] Workers entering during a recession may be more likely to have been displaced. These issues warrant special attention. To focus on spells which begin under 'normal' conditions, I therefore abstract from these spells. Focus on expansionary periods alone is warranted as in chapter 2 job polarization has been shown to occur throughout long, expansionary periods: the process of job polarization takes place during expansionary periods and is not confined to recessions or their immediate aftermath.

Survival functions are shown in figures 3.6 and 3.7. For instance, for prime aged male workers entering non-employment in the first period, the survival function at duration 1 is just above 0.5. This means that at about 50 percent of the spells which started in this period have duration of at least 1. Conversely, this implies that just under 50 percent of spells have duration less than 1. Before turning to the discussion for the changes in survival functions across periods, I first discuss some general features of these survival functions by focusing on the empirical hazard rates underlying their computation.

[24]See Jaimovich and Siu [2012].

Changes in Hazard Functions

Hazard rates for spells starting in each expansionary period are shown in figures 3.4 and 3.5 for men and women of each age group.[25] [26] First, survival functions exhibit a convex shape, i.e. they first fall sharply and then slowly. Workers generally face steeply decreasing hazard rates, i.e. workers have the highest exit probability immediately after entering non-employment, then exit probabilities decline at a decreasing rate with spell duration.

Second, survival functions are largely flat at longer durations of ten or fifteen years or more. Hazard rates indicate that for durations of ten years or longer the probability to exit is very small, although still positive. This implies that some workers exit to employment even after having been without a job for very long periods. This is suggestive that either workers decide to re-enter the labour force and exert search effort after long periods, or that workers nominally outside the labour force exhibit small but positive search effort rather than zero search effort. However, due to the data limitations these exits may also reflect workers returning to regular employment from self-employment, or returning from emigration.

Third, relative to period 1 survival functions generally shift upwards in subsequent periods, albeit not at all durations. While they generally increase at shorter durations, at longer durations they increase less strongly or, in some cases, even decrease. This pattern corresponds to shifts in hazard rates for spells starting in later periods. Hazard rates shift non-proportionally. Differences in exit probabilities for workers being non-employed for longer periods are more limited across periods. Spells starting in later periods have lower exit probabilities at shorter durations. Workers entering in later periods are less likely to leave non-employment quickly compared to workers entering in period 1. They are not much less likely to leave non-employment at longer durations. For very long durations, workers entering

[25] The empirical hazard rate corresponds to the Kaplan-Meier estimator. It is sometimes argued that Kaplan-Meier estimators assume continuous time data, whereas lifetable estimators would be appropriate for discrete time data. See, for instance, Allison [2014]. The difference is that lifetable estimators adjust for the possibility that censoring may occur before the end of the interval, in which case the number of active spells, i.e. the denominator, would be too large relative to the number of observed exits, i.e. the numerator. The adjustment consists in subtracting half the censored spells from the denominator. For simplicity, no such adjustment has been made here, as results are very similar. In any case, conclusions are not affected.

[26] Note that Stata does not report confidence intervals for Kaplan-Meier estimates of the hazard function. Generally, treatments of hazard functions focus on the continuous time hazard. A non-parametric estimate of the continuous time hazard rate can be inferred from the slope of the cumulative hazard function, for which confidence intervals are readily available. However, these estimates are said to have poor properties. I focus on the empirical hazard function as it directly relates to the survival functions discussed below. The significance of changes in survival functions is discussed later on.

non-employment in later periods are even less likely to exit to employment.[27]

Changes in Median Spell Durations: Relating Annual Transition Rates to Survival Functions

Next, I examine changes in survival functions across periods in more detail. Specifically, I discuss the extent to which changes in survival functions for spells starting after period 1 imply longer non-employment durations relative to spells starting in period 1. Unfortunately, because some spells are censored and spell duration is unobserved, one cannot readily deduce average non-employment duration from survival functions.[28]

A straightforward alternative is to instead focus on median durations, or other percentiles of the duration distribution. I focus on the 50th and 75th percentile as most spells end within the first year of observation, so lower quartiles are not always observed. The median (fourth quartile) duration corresponds to the duration at which at least 50 (75) percent of spells have ended. It can be read off the graphs as the duration at which the survival function is at or falls below 0.5 (0.75). An increase in the quartile duration implies the distribution of spell duration becomes skewed towards longer spells, i.e. workers are more likely to face longer spells. This, in turn, is suggestive that workers are non-employed for longer periods.

Before turning to the discussion, the confidence intervals shown in figures 3.6 and 3.7 indicate that changes in survival functions in later periods relative to period 1 are significant at most durations. As survival functions change non-proportionally, and increase at some durations, this generally applies to shorter and very long durations. In any case, there is evidence for significant shifts in survival functions relative to period 1. I now turn to discussing these shifts.

For male workers, median (and fourth quartile) durations generally increase for spells starting in later periods relative to spells beginning in period 1. Male workers are more likely to experience longer non-employment spells in later periods.

[27]Note that this suggests that proportional hazard models, e.g. the widely used Cox proportional hazard model, do not offer a good fit for the data when comparing survival functions across periods, as these explicitly assume that hazard rates change proportionally to a baseline hazard.

[28]Non-employment duration is treated as a random variable in survival analysis, so the average (or expected) duration is computed as the probability weighted sum over all possible durations. Because spells are censored, the maximum duration time is unobserved, however. Summing over observed durations underestimates the mean. This downward bias becomes more severe the closer periods are to the sample end, additionally making comparison across periods more difficult. Given available approaches, to compute the mean one would have to assume a continuous time parametric hazard rate, estimate its parameters using discrete time data, and then compute the closed form solution for the mean duration. See Jenkins [2005] for an exposition on deriving measures for survival time distributions.

Figure 3.4: Hazard Functions to Employment for Male Workers

(a) Young workers

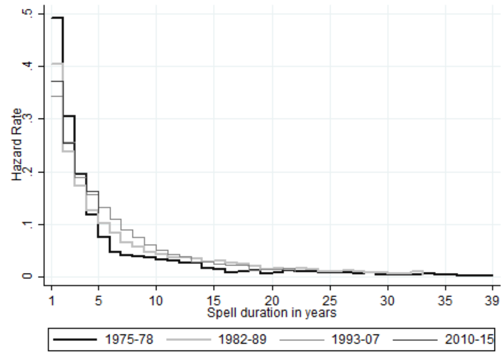

(b) Prime aged workers

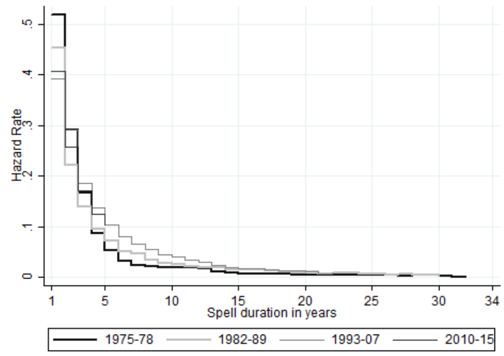

(c) Older workers

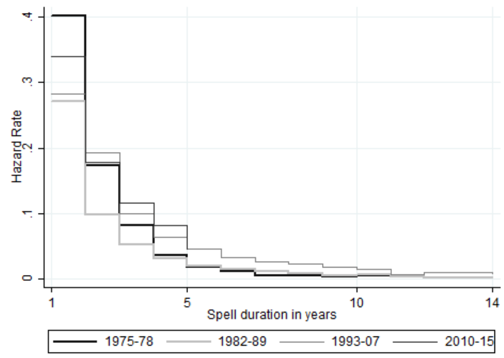

Kaplan-Meier estimates for hazard function to employment for male workers entering non-employment in expansionary periods: 1976-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.1. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure 3.5: Hazard Functions to Employment for Female Workers

(a) Young workers

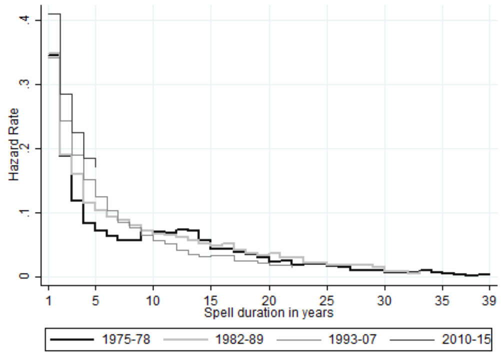

(b) Prime aged workers

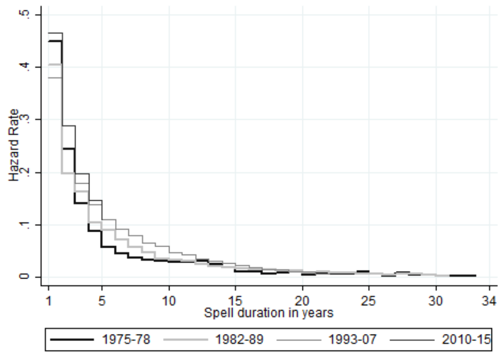

(c) Older workers

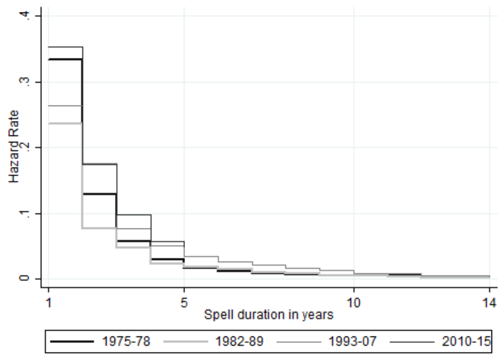

Kaplan-Meier estimates for hazard function to employment for female workers entering non-employment in expansionary periods: 1976-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.1. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

In particular, median spell duration for prime aged workers increases from one in the first period to two in all subsequent periods. The 75 percentile increases from five to as much as eleven and six in periods 3 and 5 respectively. For older male workers, median spell duration increases from two to as much as four in period 5, and to four in period 7.

The median spell duration for younger male workers remains unchanged, but one can read off changes in lower and higher percentiles from the graph. For instance, while just below 50 percent of spells end within one year in period 1, in all subsequent periods the fraction of spells ending that quickly is smaller by at least ten percentage points. The 75 percentile increases from five to eight and seven years in periods 3 and 5 respectively. Recall that young men exhibit a smaller and, using corrected rates, only temporary increase in average non-employment duration when considering corrected rates. Eyeballing survival functions for young workers suggest non-employment duration may have increased permanently, however. This may reflect that corrected exit rates understate exit rates in period 1 and 2.

For women it is more difficult to relate changes in survival functions to general increases in non-employment duration. This is because the non-monotonic change in survival functions is more pronounced for women. Clearly, young women experience declining exit rates and thus lower average non-employment duration, and this is reflected in clear downward shifts of survival functions over time. For prime aged and older female workers there is a clear upward shift of survival functions in period 3, implying longer non-employment durations. However, in period 5 we see that more spells survive at shorter durations, but less do so at longer durations, and the most recent period rather suggests that average non-employment duration decreases. In principle, this is compatible with the changes in corrected annual transition rates implying a temporary increase in average non-employment duration.

I conclude that survival functions generally suggest that workers entering in later periods were more likely to experience longer non-employment spells than in period 1. Changes in survival functions are largely compatible with longer average non-employment duration implied by exit rates from non-employment. The pattern is clearest for male workers, while conclusions for women are more difficult to draw.

Importantly, as evidence based on survival functions implies longer non-employment durations, and survival functions based on flow sampling are independent of the initial conditions problem, this assuages concerns that the decline in annual transition rates from the sample start onwards is merely an artifice of the sampling procedure.

Graphically examining survival functions serves to demonstrate a general increase in non-employment spell durations in later periods. Eyeballing these changes readily suggests non-monotonic changes in the distribution of non-employment spell durations. This suggests that job polarization may be related to heterogeneous changes in the distribution of non-employment spell durations. I discuss these changes in more detail in the subsequent sub-section.

Changes in the Distribution of Non-Employment Spell Durations

Changes in exit rates imply that job polarization is associated with workers facing, on average, longer non-employment spells in later periods. Longer average non-employment spells can reflect different scenarios. The motivation for this chapter rests on being able to distinguish whether job polarization led to longer non-employment spells as workers took more time to reallocate, or whether they failed to reallocate at all. I argue that these scenarios are related to changes in the distribution of non-employment spell durations: if workers reallocated to new jobs, but took longer in the process, the increase in average non-employment duration should reflect an increase in the fraction of longer temporary spells. If workers did not reallocate, the increase in average non-employment duration should reflect an increase in the fraction of permanent non-employment spells.

Analysing changes in survival functions allows examining changes in the fraction of temporary or permanent spells. Eyeballing hazard and survival functions, the preceding sub-sections established that changes in survival functions are compatible with longer non-employment spells, but the relationship between non-employment duration and exit probability did not change proportionally. The implication is that the distribution of non-employment spell durations changed non-monotonically. Workers in later periods are less likely to experience short spells and return to employment quickly, but they face differential changes in the probability to experience spells of various durations. To prepare the remainder of the analysis, I discuss here in some detail how the distribution of non-employment spell durations changed. In particular, I ask how the fraction of workers experiencing a temporary or permanent spell changed. These changes are informative, but may not be related to job polarization, however. This question is taken up in the following section: how much more or less likely are workers to experience temporary or permanent spells because they are less likely to leave to medium skilled jobs?

To structure the discussion, I classify spells based on their observed duration as follows: spells which end with duration less than one are referred to as short-term spells. These should correspond primarily to unemployment spells lasting less

Figure 3.6: Survival Functions to Employment for Male Workers

(a) Young workers

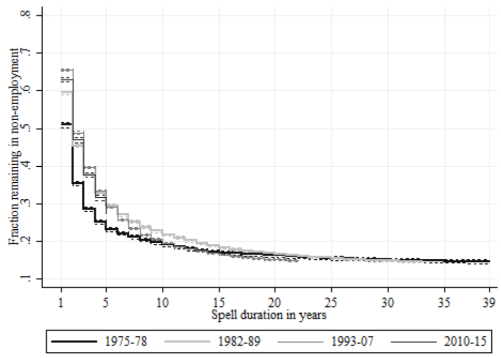

(b) Prime aged workers

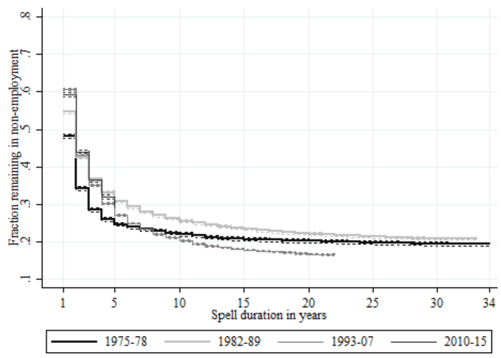

(c) Older workers

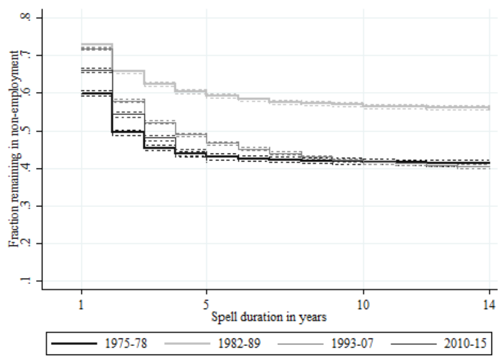

Kaplan-Meier estimates for survival functions to employment for male workers entering non-employment in expansionary periods: 1975-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.2. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure 3.7: Survival Functions to Employment for Female Workers

(a) Young workers

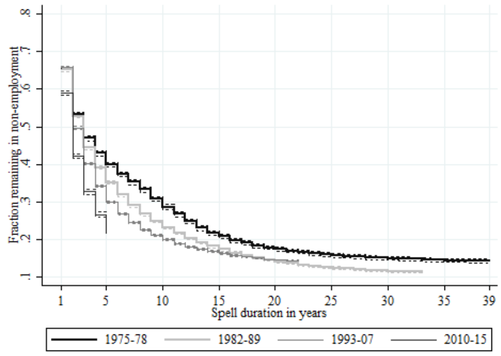

(b) Prime aged workers

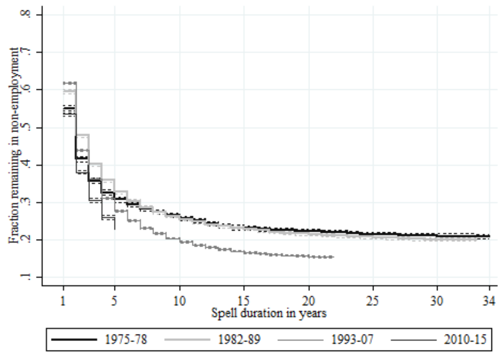

(c) Older workers

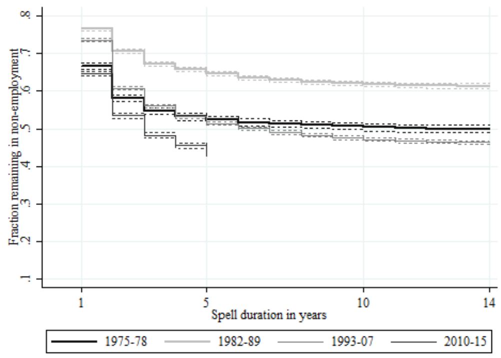

Kaplan-Meier estimates for survival functions to employment for female workers entering non-employment in expansionary periods: 1975-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.2. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

than one year.[29] Spells ending with observed durations of one to four years are classified as long-term unemployment or temporary inactivity spells. That is, spells with durations one to four should correspond to spells lasting at least one year and up to, but just short of, five years. Lastly, spells with duration five or longer are classified as permanent non-employment spells. Permanent spells proxy workers failing to reallocate. I take it that considering non-employment spells lasting at least five years is a conservative definition. It implies that workers may take more than four years searching for, and failing to find, new employment. Although it may be the case that workers continued or stopped searching for employment throughout this period, I consider them as having failed to reallocate based on the mere fact that they remained jobless over a sustained period. Either they exert no or substantially diminished search effort, the non-zero exit rate at long durations suggesting that search effort may not be zero, or they are unable to secure any job offers they would be willing to accept. In either case, I take it these workers effectively failed to reallocate to new employment upon leaving their previous job.

It is instructive for a moment to return to figures 3.6 and 3.7, which show changes in survival functions in more detail, to see how results differ if one focuses on changes in survival functions at different durations. For any given period, except at duration one, the survival function exhibits relatively even changes across subsequent durations. That is, abstracting from the fact that the survival function takes the form of a step-function, it does not exhibit large jumps. Choosing a shorter or longer duration as cut-off point for the distinction between temporary and permanent spells therefore affects results in predictable ways. Choosing a shorter duration would accentuate findings of shifts towards permanent spells. Choosing a longer duration would imply fewer permanent spells. Unless the cut-off point is chosen at much shorter or much longer durations results would remain unchanged.[30]

Before moving on to discuss these changes, I briefly describe how one can read off changes in the distribution of spell durations from changes in survival functions: using the above described terminology, a shift away from short-term spells to temporary spells is reflected by an increase in the fraction of spells surviving at duration one, while the fraction of spells surviving at duration five remains unchanged. A shift away from short-term spells to permanent spells is reflected by an increase in the fraction of spells for durations one to five by the same amount. A shift from

[29]The correspondence of this classification to actual spell lengths is not exact. This reflects the limitations discussed in the data section: only annual point-in-time observations on the employment status are made, so it is not known whether spells continue in the interval between two observations. Also, non-employment spells may reflect transitions to self-employment, emigration, or death.

[30]Below I discuss robustness checks for the main results using cutoff points of three and seven years more formally.

short-term to temporary and permanent spells is reflected in the differential increase of the fraction of spells surviving at durations one and five.[31]

What do changes in survival functions imply for changes in the distribution of spell lengths? For the ease of exposition, table 3.1 shows the value of the survival function for durations one, five and ten in period 1, as well as the percentage point change between values in period 1 and later periods. Recall that $S(1)$ gives the fraction of spells for which at least one completed year of non-employment is observed. Accordingly,$S(5)$ and $S(10)$ give the fraction of spells for which at least five or ten completed years of non-employment are observed.

Results for Male Workers Table 3.1 shows results for changes in the fraction of workers remaining non-employed at particular durations. Results for male workers are presented in the three upper panels. Male workers becoming non-employed after period 1, lasting from 1975 to 1978, generally have a much lower chance of returning to employment quickly. The fraction of spells with duration one decreases in all subsequent periods. Instead, workers experience both more temporary spells, possibly searching longer periods to find new employment, and more permanent spells, compatible with workers failing to reallocate to new jobs. The extent to which they shift to either varies substantially across periods and age groups. Workers entering non-employment in the 80s are most likely to remain permanently jobless, especially older workers. In later periods the rise in permanent spells persists, but is comparatively less important. Most workers experience temporary instead of short-term spells.

Examining changes in more detail, in period 3, lasting from 1981 to 1989, workers of all age groups are much less likely to experience short-term spells, and much more likely to become permanently jobless. It is not uncommon for these additional permanent spells to last ten years or longer. Older male workers, for

[31] Consider the following explanatory example: Suppose that $S(1)$ increased by 0.2 from 0.5 to 0.7, and $S(5)$ increased by 0.05 from 0.2 to 0.25. An increase in the fraction of spells with at least duration one corresponds to the decrease in the probability to experience short-term spell. Before the change, we have $1 - S(1) = 0.5$ short term spells. After the change, we have 1-S(1)=0.3 short-term spells. An increase in the fraction of spells with at least duration five corresponds to the increase in the probability to have a permanent spell. Before the change we have $S(5) = 0.2$, and after the change we have $S(5) = 0.25$, i.e. an increase of 0.05 percent of permanent spells. The increase in fractions with durations one to four does not itself indicate how much more likely it is to experience temporary spells, as some of the increase may correspond to permanent spells. Instead, the increase in the fraction of spells with durations one to four is given by the difference in the fraction of spells with duration of at least one minus fraction of spells with at least five. Before the change this was $S(1)-S(5) = 0.5-0.2 = 0.3$. After the change this is $S(1)-S(5) = 0.7-0.25 = 45$ percent temporary spells, i.e. a 15 pp increase. Overall, workers who enter non-employment in the later period are 20 pp less likely to experience a short-term spell, while they are 15 pp more likely to experience a temporary spell and 5 pp more likely to be permanently non-employed.

Table 3.1: Changes in Distribution of Non-Employment Spell Durations

			Level	Change to Period 1		
Gender	Age	Duration	Period 1	Period 3	Period 5	Period 7
Male	18-30	1	50.9	8.6	14.7	12.0
		5	23.2	6.3	5.8	4.6
		10	19.0	2.7	0.4	-
	31-50	1	48.1	6.5	12.6	11.2
		5	24.6	6.2	2.5	4.0
		10	21.9	3.6	-1.8	-
	51-65	1	59.9	13.0	11.9	6.2
		5	43.0	16.2	3.7	-0.9
		10	41.7	14.8	-0.2	-
Female	18-30	1	65.5	-0.3	0.3	-6.5
		5	39.9	-4.8	-10.0	-17.8
		10	28.6	-5.5	-8.7	-
	31-50	1	55.0	4.5	6.9	-1.5
		5	30.7	2.0	-3.0	-7.5
		10	25.7	-0.3	-6.3	-
	51-65	1	66.7	9.8	6.9	-2.0
		5	52.3	12.3	-0.9	-9.0
		10	50.4	11.4	-3.2	-

Change of survival functions at selected durations in later expansionary periods relative to period 1. Kaplan-Meier estimates of survival functions to employment for male and female workers entering non-employment in expansionary periods (1975-78, 1982-89, 1993-2007, 2010-15) at durations 1, 5, and 10. By age group: 18-30, 31-50, 51-65 years. Survival functions for each group and period estimated according to equation 3.2. Columns 5-7 give change in survival function relative to period 75-78. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

whom changes are most pronounced, experience an increase of about 15 pp in the probability to remain non-employed for ten years or longer, fuelled almost entirely by a decrease in the probability to experience short-term spells. For prime aged and younger workers the increase in permanent spells is more modest, but still substantial. The probability to experience permanent spells increases by about 6 pp, with half of those spells lasting ten years or longer. Again, this largely reflects a shift away from short-term spells. Evidently, the 80s constitute a period of mounting permanent joblessness. As argued below, these changes may reflect policy changes, providing workers with low employment prospects with incentives to leave the labour force and receive disability benefits.[32]

In period 5, lasting from 1993 to 2007, young and prime aged male workers experience a more pronounced fall in the probability to experience a short-term spell, a decrease of about 13 to 15 pp compared to period 1. As in period 3, young male

[32] See Center of Economic Performance [2006]; Faggio and Nickell [2005] for a discussion of these policy changes and their relation to inactivity especially of prime aged male workers.

workers are still substantially more likely to experience permanent spells, by about 6 pp, but most of the larger shift from short-term spells reflects workers experiencing temporary spells. Prime aged workers see the rise in the incidence of permanent joblessness diminished. Relative to period 1, prime aged male workers becoming non-employed in period 5 mostly experience temporary instead of short-term spells. Older male workers continue to be much less likely to return to employment quickly, by around 12 pp. Compared to period 3, in which they shifted mostly to permanent spells, in period 5 these workers are more likely to experience temporary spells. The probability to be permanently jobless remains elevated, however. It appears both prime aged and older male workers face a relative shift from permanent to temporary spells. In addition, the probability to experience very long non-employment spells decreases relative to period 1. The long period covering most of the 90s and early 2000s exhibits a moderation of the former surge in permanent joblessness. Workers continue to remain non-employed for longer periods, but they tend to return to employment in the medium-term.

In the most recent period, covering spells starting in 2010, young and prime aged male workers are still much less likely to experience a short-term spell. The probability to re-enter employment quickly remains lower by about 12 pp. Mostly they experience more temporary spells, but permanent spells remain high. They are more likely, by 4 pp, to be permanently jobless. Interestingly, older workers exhibit a tendency towards shorter spells. While still being less likely to experience a short term spell, in comparison to period 1, they are more likely to become employed after a temporary absence than in the preceding two periods, and they are even less likely to be permanently jobless than in period 1. Although the rise is no longer ubiquitous, the heightened incidence of permanent joblessness remains a concern. In line with the preceding period the shift towards temporary spells dominates.

Results for Female Workers Results for women are shown in the bottom three panels of table 3.1. Changes look very differently for women. Women do not exhibit a general tendency to be more likely to be permanently non-employed. Only older women, and to much lower extent prime aged women, entering non-employment in the 80s experienced such changes. Instead, prime aged and older women are more likely to experience temporary spells, and in successive periods exhibit a tendency towards ever shorter non-employment spells. This is in line with average non-employment duration based on corrected annual transition rates implying women only experience temporary increases in average duration.

Discussing changes by age group in some more detail, recall young women experience a decrease in average non-employment duration over the sample period. This reflects a general tendency towards shorter non-employment spells, initially fuelled by a shift from permanent to temporary spells and eventually also towards short-term spells. Young women exhibit a lower incidence of permanent joblessness in all periods after period 1. In earlier periods, the fraction of short-term spells remains largely unchanged, so young workers return to employment after temporary periods instead. In the most recent period, the fraction of short-term spells increases, so they are additionally more likely to re-enter employment quickly.

Prime aged women entering non-employment in period 3 exhibit a modest rise in the incidence of temporary spell, and to a limited extent also permanent spells, at the expense of short-term spells. However, the magnitude of these changes is small compared to other groups in this period. In period 5, prime aged women see both a decline in the incidence of short-term as well as permanent spells. As a result, workers entering during this long period face an increasingly larger probability to be without a job for longer, yet still only temporary, periods. In the most recent period, the decline in the incidence of permanent joblessness accelerates, largely shifting towards temporary spells, but prime aged women also become more likely to leave non-employment quickly.

The more modest changes in period 3 for prime aged women – a clear difference to their male counterparts - contrast with those of older female workers. The experience of older women resembles those of older men. Older women experience a sizable increase in the incidence of permanent joblessness when entering non-employment in the 80s, around 12 pp, most of which last for at least 10 years. In more recent periods, their experience is more aligned with their fellow female workers, exhibiting a tendency towards ever shorter spells. In period 5 they are more likely to experience temporary spells and less likely to be non-employed for either a short period or permanently. In period 7 they face fewer permanent and both more temporary and short-term spells.

3.3.3 Summary

To summarize, the increase in average non-employment duration for male workers is reflected by a general tendency towards longer spells in later periods. Short-term spells become less common, and the incidence of temporary or permanent joblessness increases. There is a very large rise in permanent spells during the 80s, and while the frequency of permanent joblessness remains elevated, in later periods shifts towards temporary spells dominate.

In contrast to men, for women there is no general tendency towards longer spells, in any period or overall. In period 3, when men experience a large increase in the incidence of permanent joblessness, prime aged women experience only a comparatively modest shift from short-term to temporary and permanent spells. Young women already exhibit tendency towards shorter spells. Only older women show sizeable changes comparable to male workers. However, this mixed picture in the 80s evolves into a general tendency for ever shorter spells for female workers of all age groups in periods 5 and 7.

Changes in the distribution of non-employment spell durations differ greatly across genders. The question arises whether these differences between male and female workers persist when focusing on changes driven by job polarization. For instance, do women experience shorter spells not because of differential changes in the probability to leave for medium skilled jobs, but because women are more likely to leave to other job types than men? In any case, these differences across worker groups are interesting, because they suggest that the impact of job polarization may not only be driven by aggregate factors. The overall outcome may also reflect the behaviour of workers, and the different experience of women could be seen as testimony that permanent joblessness or even longer non-employment duration in general is not inescapable. Nevertheless, results for male workers provide substantial scope for job polarization being associated with workers failing to reallocate and becoming permanently jobless. This issue, however, may have been more important in the past than in present times. It remains to be seen whether the lower rise in permanent joblessness in recent times reflects a slowing down of job polarization. The remainder of the chapter is examining to what extent these changes can be associated with job polarization, i.e. the extent to which changes in the incidence of temporary or permanent spells reflect the decline in exit rates to medium skilled employment.

3.4 Job Polarization and the Distribution of Non-Employment Durations

The preceding analysis established that the increase in average non-employment duration is associated with differential changes in the incidence of temporary and permanent non-employment spells. These changes reveal whether workers are overall more likely to be without job temporarily or permanently. To what extent are these changes associated with job polarization? The increase in average non-employment duration is associated with job polarization via the decline in the exit rate to medium skilled jobs. The remainder of the analysis examines the link between these changes and job polarization by examining the effect of changes in the hazard rate to medium skilled jobs on the incidence of temporary and permanent non-employment spells. To do so, I examine changes in survival functions in a competing risks framework, allowing to relate the incidence of non-employment durations to exits to different job types. The first, proceeding sub-section serves to introduce this framework. Subsequently, I discuss the decomposition of survival functions into changes in hazard rates to low, medium and high skilled jobs. The final sub-section discusses results.

3.4.1 Competing Risk Model

Consider the sub-sample of non-employment spells for workers who enter non-employment in period p, and who at the time they enter non-employment belong to demographic group g.

First, I discuss exits to any type of employment. Suppose that exits to employment occur in continuous time.[33] Latent exit time to overall employment is a continuous random variable, whose distribution is described by the cumulative distribution function (CDF) $F(t)$. The fact that exits occur in continuous time but are only observed at annual intervals suggests estimating the discrete-time survival function using interval-censored data.

For each spell, suppose that the time axis is partitioned into the following annual intervals:

$$(a_0, a_1], (a_1, a_2], \dots, (a_{k-1}, a_k]$$

Where $a_j = 0, 1, 2, \dots, K$ are points in time corresponding to the date of data collection, a_0 is the date the worker is first observed as non-employed and a_k is the maximum observed spell length, i.e. the latest date the spell is observed. Intervals are defined such that interval j starts the instant after a_{j-1} and includes a_j. Assuming that observations are always made with respect to the exact same reference date in each year, intervals are all of exactly one year length. It is convenient to express intervals in terms of annual duration:

$$(0, 1], (1, 2], \dots, (K-1, K]$$

The interval hazard rate for interval j is then defined as the probability to exit to employment during interval j, conditional on survival until the start of interval j:

[33] This sub-section follows the discussion of interval-censored competing risks models in Jenkins [2005].

$$\lambda_j = \text{Prob}(a_{j-1} < T \leq a_j \mid T > a_{j-1})$$

$$\lambda_j = \frac{S(a_{j-1}) - S(a_j}{S(a_{j-1})} = 1 - \frac{S(a_j)}{S(a_{j-1})}$$

The survival function for the probability of survival in non-employment until the end of interval j can be written in terms of interval hazard rates as follows:

$$S(a_j) \equiv S_j = \prod_{k=1}^{j}(1 - \lambda_k)$$

Second, consider exits to low, medium or high skilled employment, denoted as L, M and H respectively. For each spell, the latent exit time for each destination is a random variable T^A whose distribution is described by the CDF $F^A(t)$, for $A = L, M, H$. Exit to a particular destination, say low skilled employment, occurs if $T^L < T^M, T^H$. Denoting maximum observed spell length as K, exit to employment is observed if $T^A < K$ for any $A = L, M, H$, and spells are censored if $K < T^L, T^M, T^H$.

The interval hazard rate for destination A is then defined as the probability to exit to A in interval j, conditional on surviving until the start of interval j:

$$\lambda_j = \text{Prob}(a_{j-1} < T^A \leq a_j \mid T^A > a_{j-1})$$

$$\lambda_j = 1 - \frac{S^A(a_j)}{S^A(a_{j-1})}$$

Correspondingly, the survival function for the probability of survival in non-employment until the end of interval j with respect to exit to destination A is defined in terms of the destination-specific hazard rate as follows:

$$S^A(a_j) \equiv S^A_j = \prod_{k=1}^{j}(1 - \lambda^A_k)$$

It is possible to relate the overall survival function to destination-specific survival functions. As shown in section B.1 in the appendix, under the assumption of independence of competing risks, the overall hazard rate can be expressed in terms of the destination specific hazard rates:

$$1 - \lambda_j = (1 - \lambda^L_j)(1 - \lambda^M_j)(1 - \lambda^H_j) \tag{3.3}$$

Using this result, one can easily show that the overall survival function is a function of destination-specific hazard rates:

$$S_j = \prod_{k=1}^{j}(1-\lambda_k)$$

$$S_j = \prod_{k=1}^{j}(1-\lambda_k^L)\prod_{k=1}^{j}(1-\lambda_k^M)\prod_{k=1}^{j}(1-\lambda_k^H)$$

$$S_j = S_j^L S_j^M S_j^H \tag{3.4}$$

Changes in the overall survival function depend on changes in destination-specific hazard rates. Consequently, one can decompose the overall survival function into destination-specific survival functions, which themselves are functions of destination-specific hazard rates. One important caveat arises, however. Equation 3.4 assumes independence of competing risks. This assumption is worth elaborating on. Formally, it requires that random variables for destination-specific latent exit times are statistically independent. Intuitively, this means that exit times to, say, medium skill employment gives no information about latent exit times to low and high skilled employment. As a result, it is possible to estimate destination-specific hazard rates separately, treating spells to destination A as censored whenever exit to another destination occurred. Arguably this condition may not hold in reality. For instance, it is plausible that workers with high probability to exit to high skilled jobs may have a low probability to exit to low skilled jobs, as workers eligible for high skilled jobs would not accept offers for low skilled jobs. Despite its seriousness, making this assumption is common in competing risks models to avoid the substantial complications arising from dependent competing risks.[34]

It is, therefore, important to show that the above expression does not invalidate results, but provides a good approximation to overall survival functions observed in the data. Figures B.1 and B.2 show for each demographic group and period the observed overall survival function, together with the overall survival function computed according to the above equation. I refer to the latter as the predicted survival function. The overall and the destination-specific survival functions are estimated using the Kaplan-Meier estimator. Comparing observed and predicted survival functions one can see that the predicted survival function closely tracks the observed one, but that it consistently overpredicts the fraction of remaining spells by a few percentage points. This difference remains largely constant within demographic groups and over periods.[35] As the analysis examines changes in survival

[34] See, for instance, Jenkins [2005]; Allison [2014].

[35] For instance, for prime aged male workers the average deviation ranges from 3.2 pp in period

functions within demographic groups and across periods, the relatively constant upward bias in survival functions does not invalidate results, although small changes across periods should be interpreted with caution. I therefore continue with the analysis, using equation 3.4 as basis for the decomposition conducted in the next section.

3.4.2 Decomposition of Survival Functions

Based on the competing risks model introduced in the previous section, I decompose changes in survival functions into changes in hazard rates to low, medium, and high skilled jobs. This decomposition indicates how the incidence of short-term, temporary, and permanent non-employment spells changed because workers were more or less likely to leave to a particular job type.

To recall, I want to link changes in temporary or permanent non-employment spells to job polarization. The decline in the exit rate from non-employment to medium skilled jobs can account for a substantial part of the decline in medium skilled jobs, which is definitive of job polarization: job polarization occurred to large extent as workers spent longer time in non-employment instead of returning to medium skilled jobs. However, the exit rate to medium skilled jobs ignores duration dependence, and so cannot be used to analyse differential changes in the incidence of temporary or permanent non-employment spells. I turn to survival analysis instead, which explicitly models duration dependence by estimating hazard rates. The hazard rate to medium skilled employment corresponds to the duration-specific exit rate from non-employment.

Just as job polarization is related to longer average non-employment duration due declining exit rates to medium skilled jobs, I argue that job polarization is linked to changes in the distribution of spell durations resulting from changes in the hazard rate to medium skilled employment. Ultimately, the analysis indicates whether job polarization reflects workers taking longer time to reallocate to new jobs, or workers failing to reallocate and becoming permanently jobless. First, I show how one can decompose changes in the overall survival function into changes in destination-specific hazard functions. Starting from equation 3.4, taking logs, and subtracting $S(t)$ from both sides, one gets:

$$S_j = S_j^L S_j^M S_j^H$$

1, to a low of 2.5 pp in period 5, and to a maximum of 4 pp in period 7, with a standard deviation of about 0.5 pp in all cases. These numbers are representative for other demographic groups. Note that these changes are fairly small compared to the changes in survival functions discussed in the preceding section, which see changes of up to 15 pp.

$$\Delta logS_j = \Delta logS_j^L \Delta logS_j^M \Delta logS_j^H$$

$$\frac{\Delta logS_j}{S_j} \approx \frac{\Delta logS_j^L}{S_j^L} + \frac{\Delta logS_j^M}{S_j^M} + \frac{\Delta logS_j^H}{S_j^H} \tag{3.5}$$

Under the assumption of independence of competing risks, one can estimate destination-specific hazard rates separately and then compute destination-specific survival functions for each duration j as follows:

$$\hat{S}_j^A = \prod_{k=1}^{j}(1 - \hat{\lambda}_k^A) \text{ for } A = L, M, H$$

One can compute the destination-specific survival function using estimates of the destination-specific interval hazard rate λ_k^A. If no (time-varying) control variables are included, one can obtain such estimates non-parametrically using the Kaplan-Meier estimator. Alternatively, under the additional assumption that exits only occur at boundary intervals, one can estimate destination-specific interval hazard rates separately as either multinomial logit or cloglog.[36] Because my interest lies in changes in baseline hazard rates within demographic groups across periods, I construct destination-specific survival functions by estimating the Kaplan-Meier hazard rates, stratified by the period and demographic groups prevailing at the spell beginning. The advantage of the Kaplan-Meier estimates is simplicity: they require fewer assumptions and are straightforward and transparent to construct. In this case, the estimator for the interval hazard rate to destination A for interval j is:

[36] Using cloglog or multinomial logit to estimate the competing risk model allows including time varying control variables. One reason to include such variables would be to control for time fixed effects in recessionary years. If exit probabilities are substantially different in recessionary periods, one may be concerned that shorter periods, which are followed by a recession, exhibit systematically lower hazard rates. This would be of special concern for the baseline period. Longer spells which started in the first period necessarily enter the recessionary period in 1979-1981. This would, however, bias results against my findings. Also, earlier results for estimating a multinomial logit and controlling for recessionary years gave results very similar to those obtained from the more straightforward non-parametric approach. Preliminary analysis was conducted also for changes in survival functions using cloglog, and results were again very similar to results based on either logit or Kaplan-Meier estimates.

The assumption of exits occurring only at interval boundaries is necessary for multinomial logit or cloglog even if no control variables are included. Both are estimated using maximum likelihood. If exits occur within intervals the likelihood function contribution for exits to A during interval j needs to account for the probability that latent exit times to other destinations lie after the observed exit to A, but before the end of the interval. This implies the overall likelihood function cannot be separated into parts which contain only destination-specific parameters, a necessary condition to estimate destination-specific hazard rates separately. If exits only occur at boundary intervals this probability is zero, and one can estimate competing risks model by separately estimating destination-specific hazard rates. See Jenkins [2005], page 99.

$$\hat{\lambda}_j^A = \frac{h_j^A}{n_j}$$

where h_j^A denotes the number of exits to A at j, i.e. the number of spells observed at duration j as employed in A. For instance, the hazard rate for interval 1, i.e. the probability to exit non-employment during the interval starting the instant after workers enter non-employment and lasting until the end of first year, is estimated as the number of spells recorded as employed at year one divided by the total number of spells.

Tables 3.2 and 3.3 present results by age group, for male and female workers respectively. Because the approximation in equation 3.4 becomes less exact for larger percentage changes in survival functions, and because the assumption of independence of competing risks is likely to be violated, the decomposition is not exact. I also report the residual of the decomposition, i.e. the difference between the sum of changes accounted for by destination-specific survival functions and the change in the observed overall survival function. To be able to compare the contribution of changes in destination-specific hazard rates across durations I multiply each composition term by the base of the percentage change, i.e. the level of the overall survival function in period 1. The table therefore reports the contribution of changes in destination-specific hazard functions to the change in the overall survival function in percentage points. Percentage point contributions across destinations sum up to the total change. Note that, because of the approximation using logs, total changes reported in tables 3.2 and 3.3 deviate slightly from changes discussed before in table 3.1.

3.4.3 Results for Male Workers

I discuss workers by age group in turn, starting with young male workers. Following the baseline period 1, covering the late 70s, young male workers see short-term non-employment spells grow less common. Initially they are more likely to be permanently jobless instead. As time proceeds permanent non-employment spells remain elevated, but temporary spells begin to dominate the shift towards longer spells. Overall, shifts towards longer spells appear to level off in more recent periods.

Do these changes reflect job polarization? Decomposition results shown in table 3.2 imply changes are generally driven by declining hazard rates to medium skilled employment. That is, workers are more likely to remain non-employed for longer temporary spells or permanently because they are less likely to return to a medium skilled job. Associating job polarization with the declining hazard rate to medium skilled jobs, this provides evidence for job polarization leading to both longer periods of reallocation and more workers becoming permanently jobless.

Moreover, the change driven by hazard rates to medium skilled jobs becomes more pronounced over time, implying ever larger shifts towards both temporary and permanent spells. The shift towards longer spells levels off in recent periods, not because, but in spite of the sustained decrease in hazard rates to medium skilled jobs. It appears that, over time, young workers become more likely to leave to other job types. To which jobs are young workers leaving to instead? Most important are exits to low skilled jobs, which counteract the rise in temporary and permanent spells in all periods. Exits to high skilled jobs play a more mixed role. In period 3, covering spells starting during the 80s, workers are generally less likely to leave to high skilled jobs. In later periods workers evidently become more likely to exit to high skilled jobs at later durations, rendering permanent joblessness less common.

At which duration do workers become more or less likely to leave to either job type? Addressing this question provides further insight into the factors driving changes in the distribution of non-employment spell durations. For instance, if the probability to exit to, say, medium skilled jobs declines at short but increases at later durations this is compatible with job polarization leading to longer reallocation periods. Examining hazard rates allows answering this question. Figures B.3 to B.5 in the appendix show hazard rates to low, medium, and high skilled jobs for male workers. To relate changes in hazard rates to the decomposition, I show hazard rates for intervals corresponding to short-term, temporary and permanent spells (up to duration ten).[37] The interval hazard rate corresponds to the probability to leave to the respective job type during the interval, conditional on remaining non-employed until the start of this interval. For instance, the interval hazard rate for durations one to five corresponds to the probability to leave non-employment as a temporary spell, conditional on being non-employed for at least one year.[38]

Figure B.4 shows young male workers are not less likely to leave to medium skilled employment at all durations. They experience more temporary and also permanent spells because they are less likely to return to a medium skilled job during short-term and temporary spells. Fewer workers leaving at shorter durations implies more workers end up being permanently jobless. Once permanently jobless, workers

[37] I do not show hazard rates for longer spells as they do not affect decomposition results.

[38] Note that the estimation of interval hazard rates implies that longer periods generally exhibit larger exit probabilities. This is because the denominator, the number of active spells at the beginning of the interval, remains the same for all intervals, but the numerator, the number of spells exiting during the interval, generally increases with interval length. This is irrelevant for the result discussed below, however, as the comparison is generally made within each interval and across periods.

entering in periods 3 and 5, i.e. from 1982 to 1989 and 1993 to 2007 respectively, are actually more likely to return to a medium skilled job – if the permanent spell is comparatively short. Workers are more likely to be permanently jobless, but they are also more likely to return to medium skilled jobs out of permanent joblessness. This suggests that more workers considered having failed to reallocate do so eventually. One may think that this occurs as the worsening of employment prospects leads to compositional changes for permanently jobless workers: they comprise more workers who in principle seek employment, although at possibly very low search intensity. Recall that a rise in the incidence of permanent non-employment spells is interpreted as more workers failing to reallocate not because they never return to employment, but because they are without job for a sustained period. Of course these spells can also reflect workers being more likely to become self-employed and return to medium skilled jobs after some time.

Changes in the hazard rate to low skilled jobs are shown in figure B.3. Young workers are generally more likely to leave to low skilled jobs in all periods following period 1, and they also become more likely to do so over time. The probability to leave to a low skilled job increases the most at temporary durations: workers tend not to move towards low skilled jobs immediately, but after longer periods of non-employment. One may wonder whether this pattern reflects young workers being more likely to return to low skilled jobs as they were previously employed in such jobs, or whether they eventually settle for low skilled jobs after being unable to return to medium skilled jobs. I address this question later on by repeating the analysis using only spells of workers entering non-employment from medium skilled jobs.

Hazard rates to high skilled jobs, shown in figure B.5, suggest that workers are generally less likely to leave to high skilled jobs quickly. The hazard rate later on contributes to offset the rise in permanent spells, largely because workers, having been without job for longer periods, are more likely to reallocate to high skilled jobs.[39]

A similar pattern obtains for prime aged male workers. These workers experience a shift from short to longer spells, both temporary and permanent. As with

[39] Note that the hazard rate to high skilled jobs decreases slightly relative to period 1 at temporary durations in the most recent period, although the decomposition shows that changes in the survival function to high skilled jobs slightly decreases the fraction of permanent spells. This would imply that the hazard rate increases at temporary durations. This slight deviation may be explained by differences in hazard rates for longer intervals: note that the decomposition uses annual interval hazard rates. The interval hazard rate for temporary spells is computed as the fraction of all spells exiting at durations one to four, using as denominator all spells which are active at duration one. Some of these spells are censored at temporary durations, implying that annual interval hazard rates are computed over a smaller denominator.

young male workers, shifts to temporary spells gain in importance over time, while shifts to permanent spells become somewhat less pronounced.

Which of these shifts result from workers leaving less often to medium skilled jobs? Decomposition results confirm the decline in exits to medium skilled jobs can account for both the increase in longer temporary spells as well as the higher incidence of workers being permanently jobless. This is compatible with job polarization being associated with prime aged workers being permanently jobless, especially during the 80s, and taking more time to reallocate. Similar to young male workers, the push towards longer temporary and permanent spells due to declining hazard rates to medium skilled jobs grows stronger over time. There is no evidence for job polarization weakening over time. The rise in temporary and permanent spells weakens somewhat in more recent periods because workers are more likely to exit to other job types. While this is similar to the pattern observed for young male workers, one important difference obtains: in contrast to young workers, who become especially likely to leave to low skilled jobs, prime aged workers become especially likely to exit to high skilled jobs. Of course, young workers in earlier periods enter the sample of prime aged workers in later years. This may therefore suggest that workers exhibit some degree of job-upgrading over time – prime aged workers having moved from low to medium or high skilled jobs are more likely to return to such jobs – or that more seasoned workers become more reluctant to accept low skilled job offers.

What can be said about the durations at which workers are more or less likely to exit to low, medium, and high skilled jobs? Figure B.4 in the appendix reveals prime aged male workers are generally and increasingly less likely to exit non-employment quickly by returning to a medium skilled job. This explains why for prime aged workers it is less common to experience short-term spells. The sustained rise in the incidence of permanent joblessness reflects a more complex pattern: the probability to leave to medium skilled employment at a temporary duration decreases in periods 3 and 7, but returns close to its initial level in period 5. What does this imply? The higher incidence of permanently jobless workers in period 5 is compatible with workers having the same probability to leave a temporary spell as in period 1. It is higher merely because fewer workers leave quickly. The discontinuous change in the hazard rate for temporary spells is itself of interest, as it points to varying chances of reallocating to medium skilled jobs after longer periods, with workers being comparatively successful in period 5. The question arises what was different for these workers to be more successful? At any rate, later periods see workers who are permanently jobless being more likely to exit to medium skilled jobs. Again, the rise in the fraction of permanently jobless workers is accompanied by workers being more likely to return from permanent joblessness to a medium skilled job.

Figure B.5 in the appendix shows hazard rates to high skilled jobs remain fairly constant in period 3, exhibiting a general rise only in later periods. From the 90s onwards, prime aged male workers are more likely to exit to high skilled jobs at all durations. This general increase may reflect changes in educational attainment. If so, assuming educational attainment increases at an early age, it is interesting to note that younger workers in later periods do not see a similar, general increase. One might think that educational attainment increases at an early age, most commonly at the age of twenty one. The fact that hazard rates to high skilled jobs increase in later generations for prime aged, but not for young workers, seems puzzling on this account. All the same, the impact of job polarization on workers suffering longer non-employment spells is diminished to some extent as these workers return to high skilled jobs after short-term, temporary, as well as permanent non-employment spells. The hazard rate to low skilled jobs similarly exhibits a general increase in the probability to leave to low skilled jobs. The increase is largest at longer durations, and grows over time. In absolute terms, however, the change is small compared to the hazard function to high skilled jobs, thus the smaller impact on the shift towards temporary and permanent spells.

Finally, I discuss older male workers. Older workers experience a general shift away from short-term spells. In period 3, from 1982 to 1989, this shift implies workers are much more likely to be permanently jobless. However, this initially large shift to permanent spells vanishes in subsequent periods, leading instead to more temporary spells. Ultimately, permanent joblessness is less common for workers entering in period 7, from 2010 to 2015, than in period 1, from 1976 to 1978.

Table 3.2 shows these changes are again driven by changes in the probability to leave to medium skilled jobs, compatible with job polarization encompassing the surge in permanent joblessness in the 80s, and the subsequent rise in workers taking longer time to reallocate. In particular, declining hazard rates to medium skilled employment imply an increase in the share of temporary spells and permanent joblessness in all periods. In contrast to young and prime aged male workers, the magnitude of this general decline grows weaker over time. This pattern is generally reinforced by changes in exits to other job types: in period 3 workers are less likely to exit to low and high skilled jobs, adding to the rise in permanent joblessness. Afterwards the rise in temporary spells is partly offset, and permanent joblessness is less common, because older workers leave low or high skilled jobs more frequently. The tendency to leave to high skilled rather than low skilled jobs as one gets older continues, with exits to high skilled jobs playing a more pronounced role than exits to low skilled jobs for older workers. At which durations do older workers become more or less likely to exit to a particular job type? Older workers are generally less likely to return to medium skilled jobs quickly. The decline in the hazard rate at a low duration persists throughout all periods. Later on, the rise in temporary spells increases, and reverses eventually for permanent spells, because they become more likely to return to a medium skilled job during a temporary spell. Just as younger and prime aged workers they are also more likely to exit to medium skilled jobs from permanent joblessness. They generally become more likely to leave to low or high skilled jobs. Older workers increasingly leave non-employment quickly to high skilled jobs, while they leave to low skilled jobs more reluctantly: They are less likely to take on low skilled at a short duration, but more likely after spending some time non-employed.

3.4.4 Results for Female Workers

Results are generally very different for female workers. Consider, for example, young women. Young women exhibit a general trend towards ever shorter spells, ranging over the entire sample period.

Decomposition results are shown in table 3.3. Young women experience ever shorter spells because they become more likely to exit to any type of job. A rise in exits to medium skilled jobs importantly contributes to this tendency, but only by reducing the frequency of permanent spells. Young women experience fewer short-term spells because of a decline, at apparently shorter durations, in exits to medium skilled jobs. It is worth noting that the implied reduction in short-term spells is much smaller compared to any male worker group, and it is also smaller than the implied reduction in permanent spells. Changes in exits to medium skilled jobs increase the frequency of temporary spells. Overall, women return to employment more quickly mostly because they exit to low skilled jobs. The exceptional pattern observed for young women is thus driven by exits to other job types, with medium skilled jobs playing an important but mixed part. Job polarization may have contributed to longer spells as young women take more time to reallocate. Despite the presumed decline in demand for medium skilled jobs, however, young women also manage to experience less permanent joblessness – in absolute terms – by leaving to medium skilled jobs. In contrast to other workers young women mainly spend more time to reallocate, and are comparatively successful at that. If job polarization implied a general worsening of labour market prospects for workers seeking medium skilled jobs, the question arises what is different about young women?

Looking at hazard rates in figure B.7 in the appendix confirms that women are less likely to leave to medium skilled jobs at short durations, but they become more likely to leave at longer durations. This would be compatible with job polarization leading to longer reallocation periods. The push towards shorter spells, from the 90s onwards, reflects a general increase in the exit probability to high skilled jobs, occurring at all durations. This may be taken to suggest that the increase reflects increasing educational attainment, assuming that more educated workers generally have a higher probability to exit to high skilled jobs. The reduction in spell lengths due to exits to low skilled jobs mainly reflects young women being more likely to exit to such jobs at longer durations.

I turn next to prime aged female workers. After initially experiencing modestly longer spells, these workers spend less and less time non-employed in subsequent periods.

The decomposition shows this pattern is not generally driven by changes in hazard rates to medium skilled jobs. The initial push towards longer spells reflects prime aged women being less likely to exit to low and medium skilled jobs. Especially medium skilled jobs contribute to a modest increase in permanent spells, but these effects seem generally small. In later periods, as with younger female workers, short-term spells become less common for prime aged women because of fewer exits to medium skilled jobs, but so does permanent joblessness. A similar pattern applies with regard to exits to low skilled jobs. There is a clear picture of prime aged women experiencing shorter spells because they exit to high skilled jobs more often. From the 80s onwards, this reduces both the incidence of permanent as well as temporary non-employment spells. Again, results are compatible with job polarization having a limited impact, leading prime aged women to take more time to reallocate, but not to more permanent joblessness.

How exactly do hazard rates change? In line with the above pattern, the hazard rate to high skilled jobs, shown in figure B.8, confirms that prime aged women are generally more likely to exit to high skilled jobs at all durations. Figure B.6 shows prime aged women are less likely to experience short-term spells because they are less likely to exit to low and medium skilled jobs at short durations. As young women, they are more likely, however, to exit to such jobs after being non-employed for some time. Again, this can be seen as evidence of job polarization impacting on women by prolonging the time they take to reallocate to a new job, without leading to permanent joblessness.

Finally, I turn to discussing results for older female workers. Older women entering non-employment in the 80s face a large rise in permanent joblessness, and fewer short-term spells. They subsequently experience shifts towards ever shorter spells.

The decomposition shows that the overall pattern cannot be explained in terms of changes in hazard rates to medium skilled employment alone. The surge in permanent joblessness in period 3 is driven by older women leaving to both low and medium skilled jobs less often. The following push towards ever shorter spells occurs, on the one hand side, as older workers become generally more likely to exit to high skilled jobs. On the other hand side, the decline in exits to low and medium skilled jobs first weakens, and then reverses. Fewer exits to medium skilled jobs always contribute to older women being less likely to exit non-employment quickly, but permanent joblessness ultimately decreases because older women become more likely to leave to medium skilled jobs, although apparently only at later durations. Again, evidence suggests that job polarization may be associated with a modest shift towards longer periods required for reallocation, but not permanent joblessness.

How do changes in hazard rates relate to this pattern? Not surprisingly, initially the hazard rate to medium skilled jobs declines at short-term and temporary durations. Later on, older women are less likely to experience permanent joblessness because they are more likely to exit to medium skilled jobs at temporary durations, and they are less likely to return to employment quickly because the hazard rate to medium skilled jobs declines at short-term durations. Exits to high skilled jobs remain fairly constant at first, and subsequently increase at all durations, thus exhibiting a general upward shift: older women are generally more likely to exit to high skilled jobs. Changes are similar for exits to low skilled jobs. The initial decline is more pronounced, and the subsequent increase occurs only for longer durations, however. Older workers remain less likely to leave to low skilled jobs at shorter durations.

3.4.5 Results Conditional on Previous Medium Skilled Employment

The previous section discusses results not conditional on previous employment. It is not possible distinguishing whether changes in hazard rates to, say low skilled jobs, only apply to workers who enter from low skilled jobs, or whether they reflect genuine reallocation, i.e. workers who previously worked in a medium skilled job become more likely to reallocate to low skilled jobs. Put differently, do the above results hold for all workers, or would they differ markedly for workers who enter non-employment from medium skilled jobs.

To address these questions, I repeat the above decomposition for a subset of spells: I restrict the analysis to spells of workers who, before entering non-employment, worked in medium skilled jobs. Such workers becoming more or less likely to exit to low or high skilled jobs provides clear evidence of workers reallocating to other employment types out of non-employment. Additionally, conducting the analysis conditional on previous employment goes some way towards controlling for possible compositional differences across genders. Women may experience smaller shifts towards longer spells because of declining exits to medium skilled jobs as, in earlier periods, they are already less likely to work in medium skilled jobs.

Results are shown in the appendix in tables B.3 and B.4. Results are by and large very similar to those obtained for unconditional spells. Workers entering from medium skilled jobs do not face widely different changes in the incidence of temporary or permanent spells compared to workers entering from any type of employment, including low and high skilled jobs.

Initial changes in the probability to leave to medium skilled jobs imply somewhat larger shifts towards temporary as well as permanent spells. Workers who enter non-employment from a medium skilled job are more likely to be non-employed for longer temporary periods, or to be permanently jobless. This would imply that it is especially medium skilled workers who take longer to reallocate, and who fail reallocating. This is the case for male workers but, especially, for prime aged and older women. Focusing on changes experienced by medium skilled workers implies that prime aged and older women face changes more similar to their male counterparts. Changes implied by fewer exits to medium skilled jobs in period 3 in fact suggest women and men of the respective age group faced changes of the same magnitude. This may suggest that some of the different experiences encountered across genders reflect compositional differences: arguably, women are less likely than men to work in medium skilled jobs in earlier periods, and workers previously in medium skilled jobs experience larger shifts because of the decline in exits to medium skilled jobs. Results for the conditional analysis demonstrate, however, that the differences remain in later periods. Interestingly, young women entering non-employment from medium skilled jobs actually experience fewer permanent non-employment spells because of more frequent exits to medium skilled jobs at later durations. Conditioning on previous medium skilled employment cannot explain away differences across genders. In any case, differences between conditional and unconditional results vanish in later periods. Whatever distinguishes these workers from those entering from low or high skilled jobs disappears in more recent periods.

Small differences obtain. In particular, results conditional on previous employment in medium skilled jobs hint at a larger role played by exits to high skilled jobs in reducing non-employment spell lengths. The fact that results are largely similar, however, suggests that previous results reflect genuine upgrading. Workers reallocate during non-employment from medium skilled to 'lovely' jobs, although the extent of this happening is comparatively small. Recall that most workers reallocate to other medium skilled jobs. That these results are somewhat more pronounced may suggest that medium skilled workers are somewhat more likely to upgrade. This applies to male and female workers.

Results for exits to low skilled jobs are also quite similar for conditional and unconditional spells, although a bit less pronounced. This may be taken to suggest that workers entering non-employment from medium skilled jobs are somewhat less likely to downgrade to a low skilled job. This is compatible with the stronger results for the unrestricted sample reflecting workers becoming more likely to work in low skilled jobs, and also more likely to return to a low skilled job when becoming non-employed. Again, this holds across genders.

3.4.6 Results for Shorter and Longer Cutoff Durations

There are no definitive criteria for choosing the cutoff duration to distinguish longer temporary from permanent spells. Intuitively, a permanent spell should be sufficiently long to warrant that the period of joblessness itself implies a failure to reallocate. Alternatively, one may think of permanent spells as spells sufficiently long to imply that workers have been inactive for at least some time during the spell. I choose a cutoff duration of five years as I take this to be a conservative measure for permanent spells on both accounts.[40] To show that conclusions do not depend on a particular cutoff duration, tables B.1 and B.2 give results for the decomposition at durations three and seven, in addition to durations one, five, and ten focused on in the main analysis. Results shown in the tables show that qualitative conclusions are generally unaffected by the choice of the cutoff point. If one were to choose a shorter duration of three years to distinguish permanent from temporary spells, one would generally conclude larger shifts towards non-employment spells, driven by the same pattern of job type-specific hazard rate changes. If instead one were to choose a longer duration of seven years, one would still observe sizable shifts towards permanent spells, again driven by the same hazard rate changes, but of smaller magnitude. In a few instances some additional results emerge which, however, do not affect general conclusions. For instance, for prime aged workers one can observe additional shifts within temporary spells: the fraction of spells with duration three increases by more than the fraction of spells with duration one, implying that spells shifted

[40] I chose duration one as cutoff for short-term spells as one year is the shortest observed duration.

from short-term and shorter temporary spells towards longer temporary spells.

3.4.7 Summary

A general pattern emerges: in all periods following period 1, and for all groups, workers are less likely to return to employment quickly because of a decline in the probability to take up a medium skilled job at a short duration. This is compatible with job polarization leading to longer spells as workers take more time to reallocate.

Job polarization can also be associated with workers becoming permanently jobless. However, increased incidences of permanent spells vary substantially across groups. Men are generally more likely to become permanently jobless because of lower exit rates to medium skilled jobs also at longer durations. However, women do not experience a general increase, and even experience fewer permanent spells because they leave to medium skilled jobs more often. These different experiences may reflect differences in composition, which are not accounted for, or behaviourial differences. Conducting the analysis conditional on previous employment goes some way towards controlling for compositional differences. While prime aged and older women entering non-employment from medium skilled jobs do appear to be more similar to their male counterparts, large differences remain, at least in later periods. A further difference across genders, contributing to the diverging tendencies across men and women, is the more pronounced increase in exits to other job types experienced by women. Women appear to be more successful at reallocating to alternative job types. However, women also experience a less pronounced, or even reversed decline in hazard rates to medium skilled jobs, so women also appear to be more successful at reallocating to other medium skilled jobs. This difference across genders cannot be explained solely in terms of women being more willing or able to take on low and high skilled jobs than men. Differences with regard to exits to medium skilled jobs remain important. The implication is that job polarization is associated with permanent non-employment spells, but permanent joblessness is not inescapable. If these differences are the result of women behaving differently, an important question for future research is to examine how their behaviour differs.

The rise in permanent joblessness generally seems stronger in earlier periods. However, as the analysis has shown, the subsequent weakening of the rise, and differences across gender for early periods, reflect workers reallocating to other job types. It stands to reason that for many workers successfully reallocating to new a job requires considering other job types. Incidentally, hazard rates to medium skilled jobs generally reveal that, despite their decrease after period 1, the probability to exit to medium skilled jobs remains high compared to exits to low and high skilled jobs, also at longer durations. Despite the decline in the hazard rate to medium skilled jobs, in absolute terms many workers take more time to reallocate, or fail to reallocate and remain without job for a sustained period, but eventually return to a medium skilled job. These patterns largely apply to workers entering non-employment from all job types, as well as to workers entering from medium skilled jobs only. It appears that some of the increase in the probability to exit to low skilled jobs reflects workers already working in such jobs returning to similar jobs. However, a substantial part of the rise in exits to low and high skilled jobs reflects genuine reallocation. Workers who are previously employed in a medium skilled job take up other job types, often after long periods of joblessness.

Interestingly, the increase in very long spells, which I refer to as permanent, is accompanied by a higher probability to return to medium skilled jobs at very long durations. More workers are non-employed for long periods, but they also return to medium skilled jobs more often. This may reflect a change in the composition of the permanently non-employed: the additional workers are unable to return to employment for a sustained period, but they remain attached to the labour force in principle, possibly exhibiting low but non-zero search effort.

Additional interesting issues emerge. Both the hazard function to low and high skilled jobs generally shift upwards. However, the probability to exit to a low skilled job increases primarily at longer durations, and often declines at a short duration. Does this suggest workers accept low skilled jobs only reluctantly, after they learn about the unavailability of better jobs? Future research may shed light on this question. With regard to exits to high skilled jobs, one generally observes an upward shift at all durations. Assuming that more highly educated workers are generally more likely to exit to high skilled jobs, at all durations, this upward shift may reflect increasing educational attainment. Clearly, the question how results would differ by educational attainment is an important one.

Some apparently systematic differences across age groups occur. Young workers generally become more likely to exit to low skilled jobs, while prime aged and older workers become relatively more likely to exit to high skilled jobs. As any increase in exit rates to high skilled jobs reflecting educational attainment should apply to all groups, and in fact to older age groups only with a lag, one may wonder whether these differences reflect job upgrading, or workers of a certain age being more reluctant to take up low skilled jobs. Interestingly, if it is the case that they are more reluctant, older workers do comparatively well regardless. Despite the comparatively smaller increase in the probability to leave to low skilled jobs, they generally experience large declines in the share of permanent non-employment spells.

Table 3.2: Decomposition of Survival Functions for Male Workers

Young Workers

		Total	Contribution			
Period	Duration	Change	Low	Med	High	Residual
	1	8.0	-0.3	6.0	0.8	1.5
3	5	5.6	-0.7	4.8	0.6	0.9
	10	2.6	-0.9	2.5	0.3	0.7
	1	12.9	-0.8	10.7	0.9	2.1
5	5	5.2	-1.7	6.3	-0.2	0.8
	10	0.4	-2.1	2.9	-0.9	0.5
	1	10.8	-2.5	11.0	0.9	1.4
7	5	4.2	-3.5	7.4	-0.1	0.4

Prime Aged Workers

		Total	Contribution			
Period	Duration	Change	Low	Med	High	Residual
	1	6.1	-0.1	4.7	0.2	1.3
3	5	5.5	-0.1	4.4	0.3	0.9
	10	3.3	-0.3	2.8	0.0	0.8
	1	11.2	0.2	8.8	0.0	2.3
5	5	2.4	-0.5	3.6	-1.5	0.8
	10	-1.8	-0.9	0.7	-2.3	0.6
	1	10.0	-0.3	9.4	-0.8	1.7
7	5	3.7	-1.0	6.3	-2.2	0.7

Older Workers

		Total	Contribution			
Period	Duration	Change	Low	Med	High	Residual
	1	11.8	0.8	8.6	0.7	1.7
3	5	13.7	1.0	10.1	1.1	1.5
	10	12.7	0.8	9.4	1.0	1.4
	1	10.9	0.5	9.1	-0.2	1.4
5	5	3.5	-0.2	4.6	-1.6	0.8
	10	-0.2	-0.7	2.0	-2.2	0.7
	1	5.9	0.2	7.0	-1.7	0.4
7	5	-0.9	-0.7	3.3	-3.5	0.0

Decomposition of changes in survival function relative to period 1975-78 at durations 1, 5, and 10. Decomposition based on equation 3.5. Survival functions based on Kaplan-Meier estimates for male workers entering non-employment in expansionary periods (1975-78, 1982-89, 1993-2007, 2010-15). By age group: 18-30, 31-50, 51-65 years. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Table 3.3: Decomposition of Survival Functions for Female Workers

Young Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
	1	-0.3	-0.4	0.1	0.1	-0.1
3	5	-5.1	-2.6	-1.9	-0.2	-0.4
	10	-6.1	-3.4	-2.2	-0.1	-0.4
	1	0.3	-0.8	2.0	-0.7	-0.2
5	5	-11.5	-4.1	-3.6	-2.5	-1.3
	10	-10.4	-3.9	-3.2	-2.3	-1.0
	1	-6.9	-4.9	1.5	-1.4	-2.2
7	5	-23.6	-11.9	-4.9	-3.2	-3.5

Prime Aged Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
	1	4.3	1.4	2.0	-0.3	1.2
3	5	1.9	0.8	1.1	-0.7	0.7
	10	-0.3	-0.3	0.2	-0.8	0.6
	1	6.5	2.8	2.8	-0.8	1.7
5	5	-3.1	0.3	-1.4	-2.3	0.2
	10	-7.2	-1.7	-3.0	-2.6	0.0
	1	-1.5	1.0	0.4	-2.2	-0.8
7	5	-8.6	-1.7	-2.1	-3.4	-1.4

Older Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
	1	9.1	2.9	4.3	0.3	1.6
3	5	11.0	3.7	5.4	0.4	1.5
	10	10.3	3.1	5.4	0.3	1.4
	1	6.6	2.3	3.7	-0.6	1.2
5	5	-0.9	0.0	0.4	-1.8	0.5
	10	-3.3	-1.4	-0.4	-2.1	0.5
	1	-2.0	-0.5	0.6	-1.4	-0.7
7	5	-9.8	-3.7	-2.5	-2.6	-1.0

Decomposition of changes in survival function relative to period 1975-78 at durations 1, 5, and 10. Decomposition based on 3.5. Survival functions based on Kaplan-Meier estimates for female workers entering non-employment in expansionary periods (1975-78, 1982-89, 2010-15). By age group: 18-30, 31-50, 51-65 years. Columns 4-7 give contribution of changes in hazard functions to low, medium, high skilled employment and decomposition residual. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

3.5 Discussion

This chapter set out to examine the link between job polarization and the duration of non-employment spells. Specifically, I asked whether the decline in the non-employment outflow rate to medium skilled jobs, which is associated with job polarization in the UK, is the result of all workers taking more time to reallocate to new jobs, or whether it reflects some workers being permanently jobless? To link changes in non-employment durations to job polarization, I identified changes in the distribution of spell durations resulting from changes in hazard rates to medium skilled jobs. As annual transition rates from non-employment to medium skilled jobs can account for large parts of job polarization, this is suggestive for polarization being associated with these distributional changes.

I interpret a spell lasting at least one but less than five years as a longer temporary spell, indicating that workers take more time to reallocate. I referred to non-employment spells lasting at least five years as a permanent, and I take them to imply a failure to reallocate. A general rise in longer, temporary spells indicates that the decline in the non-employment outflow rate results from all workers taking more time to reallocate. A rise in the incidence of permanent spells implies that the outflow rate drops because some workers are disproportionately affected. Polarization would indeed be very costly for the afflicted worker: the long duration of non-employment gives rise to the view that these workers failed to reallocate. The basic reasoning, that the destruction of redundant jobs does not result in non-employment because it gives way to new, possibly better employment opportunities, would need to be qualified. Workers may return to employment eventually, but they do so only after protracted periods of joblessness.

What do results suggest for the impact of job polarization on non-employment spell durations? The decline in outflow rates to medium skilled jobs results in both longer temporary spells as well as workers being more likely to be permanently jobless. The rise in permanent spells is small compared to the increase in temporary spells, but it is often large in absolute terms, and generally persistent for male workers. Results imply that job polarization leads to a general shift from short-term to longer temporary spells as well as permanent joblessness. Job polarization is associated with workers generally taking more time to reallocate, and some worker groups are disproportionately affected as they fail to reallocate for very long periods.

Some additional patterns emerge. Distributional changes for non-employment duration spells largely occur at durations of up to ten years. There is no general increase in the number of spells lasting longer. In fact, the shift towards longer spells is accompanied by rising hazard rates at longer durations, also to medium skilled jobs. Workers are more likely to be jobless for very long periods, i.e. to experience what I refer to as permanent joblessness, but they return to employment eventually. This holds true for all workers except for older workers becoming non-employed during the 80s, who indeed appear to never return to employment. Also, shifts in the distribution of spell durations towards longer temporary or permanent spells vary over time and across groups. There is a sustained shift from short-term towards longer temporary and permanent spells for male workers. Women, however, are generally less severely affected in early periods, in that they experience smaller shifts towards longer spells, and in later periods they generally face shorter spells. Young women even face shorter spells because they leave to medium skilled jobs more often. This suggests that the rise in longer non-employment durations may not reflect aggregate factors alone. Instead, the observed patterns result from the interaction of aggregate changes, such as the decline in job creation for medium skilled jobs, with compositional variation across groups, e.g. in educational attainment, or with different behavioural responses. Given the stark difference between changes observed for male and female workers, an important question to be addressed by future research is whether male workers could avoid permanent joblessness if they changed their characteristics or behaviour.

Of course, these findings have to be seen in light of their limitations. I do not observe search effort, which would allow classifying workers into those seeking and not seeking employment, and neither do I observe the cause for them entering non-employment, e.g. whether they are dismissed for personal reasons or are being made redundant. A further limitation is that non-employment may be mismeasured. It may reflect workers becoming self-employed, emigrating, or dying before the age of 65. My analysis can therefore only be suggestive about substantive interpretations of the observed patterns. It is nevertheless instructive to discuss possible implications of my results in some more detail in light of these limitations.

In particular, it is instructive to relate findings to a standard general equilibrium model featuring job search, search effort and a participation decision, and to ask whether such a model is consistent with the observed patterns. The relevant seminal model here is introduced in Pissarides [2000].[41]

To start the discussion, one may think of job polarization as a decrease in the productivity of medium skilled job matches relative to the fixed income of unemployed workers.[42] Such a decrease is plausible if one thinks of job polarization

[41]In particular, I refer to the baseline model set out in chapter 1 and extended in chapters 5 and 7 for search intensity and participation decisions.

[42]It is well known that a decrease in productivity decreases wages, but does not affect labour

as resulting from accumulating capital which substitutes for medium skilled jobs, or routine cognitive tasks, while possibly complementing low and high skilled jobs, or non-routine tasks.[43] In this case, the net match surplus of medium skilled jobs, and firms' profits, decline, thus decreasing job creation of medium skilled jobs. Wages fall, but not enough to compensate for the drop in profits. In steady state, job polarization implies that demand for medium skilled jobs falls relative to supply, leading to a decline in wages and labour market tightness.[44]

What is the effect of such a scenario on unemployment duration, and labour force participation? Consider first unemployment duration. The fall in labour market tightness means that job finding rates are lower. As expected unemployment duration is just the inverse of the job finding rate, this implies that average unemployment duration increases. This is further amplified if one accounts for search intensity due to a discouraged worker effect. Search intensity in the model can be seen as the average fraction of time workers spend on finding employment, acting to increase the effective number of unemployed workers looking for employment. Workers decide how intensively to search by trading the costs, which are convex in search effort, against the gain in expected net worth resulting from a higher job finding probability. Because labour market tightness and wages are both lower, job polarization acts to reduce this gain for medium skilled workers. The implication is that they reduce search effort in equilibrium. As workers enter the market less frequently to look for jobs, average unemployment duration increases even more.[45]

Second, what is the effect of job polarization on the participation rate? Workers decide to enter the labour market and start searching for a job if they are better off searching than remaining inactive. It is easiest to think of income, or utility, gained during inactivity as given exogenously and varying across workers. The participation decision then depends on the value of unemployment, which reflects the fact that search is costly but yields an expected increase in net worth from finding employment. Workers become unemployed and search for jobs if the value of un-

market tightness in this model, if the productivity decline is fully absorbed by wage adjustments. For the wage to not adjust fully, it is necessary that the worker's threat point in wage negotiations increases relative to the match surplus. This is the case if fixed income of unemployed workers decreases less than the match surplus.

[43] Cortes et al. [2016] make a corresponding argument in the context of a frictionless model.

[44] In the basic model labour supply is exogenous. In a model with participation decision, labour supply may responds to changes in productivity. However, the adjustment of the participation rate does not affect the drop in labour market tightness and wages as the model is independent of the scale of the economy. The participation rate is solved for recursively as a function of equilibrium labour market tightness.

[45] This refers to a symmetric equilibrium in which all workers choose the same search intensity and no worker has an incentive to deviate. In principle, a decline in aggregate search intensity increases the job finding probability for a single worker because of congestion externalities.

employment exceeds their reservation utility. A decline in labour market tightness and wages both renders unemployment less attractive: it takes longer to find a job, and jobs are worth less. Labour force participation decreases as a result, and more workers become inactive.[46]

If one is prepared to think of job polarization as a permanent decline in relative productivity for medium skilled jobs, the standard search and matching model suggests a permanent rise in unemployment duration, driven by a decline in job creation and search effort, as well as an increase in inactivity. The general shift from short-term towards longer temporary spells, driven by changes in the hazard rate to medium skilled jobs, is consistent with the increase in average unemployment duration. Also, it seems plausible that workers who are non-employed for very long periods are inactive at some point during their spell. The model's implications are then also consistent with the sustained rise in permanent joblessness experienced by male workers.

What does the model imply about possible explanations for the variation of the observed patterns across groups and over time? Consider first differences in the magnitude of the rise in temporary spells across groups.

The basic model implies that unemployment duration depends on labour market tightness and search effort. The general increase in temporary spells might reflect a fall in labour market tightness. Differences in the magnitude of this rise in temporary spells across groups could point to compositional differences affecting average search effort. Suppose, for instance, that women experience larger gains in educational attainment than men. Assuming that higher educated workers have a larger gain in net worth associated with finding employment, women's search effort would increase relative to men. Recall in this context that results already indicate that prime aged and older women experience shifts more similar to their male counterparts if conditioning the analysis on medium skilled workers.

What could account for differences in the increase in permanent joblessness across groups? The model suggests that participation rates depend on the value of unemployment and reservation utility. The fact that women experience more limited shifts towards permanent joblessness could reflect differences in the gain from job search, or differences in reservation utility. Again, compositional changes could account for the more limited increase in permanent joblessness if women differ from men in possessing characteristics making them more likely to find relative well-paying employment, e.g. being better educated. To account for the observed

[46] Note that a decline in productivity also decreases the equilibrium participation rate if it decreases in proportion to the fixed income of unemployed workers, because the costs of hiring in the model depend on productivity.

patterns in terms of variation in reservation utility, women would need to have generally lower reservation utility than men, which does not seem readily plausible. In fact, explaining differences in shifts towards permanent spells, or inactivity, in terms of men having higher reservation wages would suggest that permanently non-employed male workers are better off, in absolute terms, than employed women with similar characteristics: they choose not to work because they enjoy higher utility being inactive. This seems counterintuitive, especially for permanently jobless prime aged male workers. A related explanation could be that preferences differ among men and women but are unrelated to reservation utility. Anecdotal evidence suggests that male workers may find it more difficult to adopt to changing circumstances by switching occupations, or moving to areas with better employment prospects. In terms of the model, male workers previously holding medium skilled jobs may experience larger disutility from accepting available jobs which are different from their previous occupation.[47]

Finally, what can explain the variation over time in shifts towards temporary and permanent spells? To account for these changes in terms of the model, either the gain from job search has to increase over time, or reservation utility has to fall. Recall that women differ from male workers as they experience shifts towards ever shorter spells in recent years. Again, compositional changes affecting the gain from job search could account for this pattern, for instance if women continue to outpace male workers in terms of educational attainment. If men also increase their educational attainment, one would expect less workers to become permanently jobless in later years.[48] An alternative explanation is related to policy changes. It has been noted that, beginning in the mid 1970s, inactivity rose substantially among men, and that this rise was associated with a surge in claimants for invalidity

[47]Recall that results above demonstrate that shifts to other than medium skilled job types generally play a larger role for female than for male workers. Also, note in this context that shifts towards inactivity appear to be concentrated among low skilled prime aged workers (Faggio and Nickell [2005]; Center of Economic Performance [2006]).

[48]A related but prima facie less convincing explanation could be that workers change their search behaviour, and revise their decision to accept jobs or join the labour market, as they learn about their true employment prospects. Cahuc et al. [2014] discuss a basic one-sided search model in which workers receive job offers from a known wage offer distribution and choose a reservation wage to accept or decline offers. Job polarization would reflect a decline in the job offer arrival rate as well as a leftward shift in the wage offer distribution for medium skilled jobs. If workers have knowledge about the initial wage offer distribution, but need to learn about the shift implied by job polarization, their reservation wage will be too high at first. They reject offers they perceive to be unattractive, expecting better offers to arrive soon. As these offers do not occur, they revise their expectations and lower their reservation wage, becoming more likely again to accept job offers. Still, one may wonder whether it takes several years to revise expectations, and while this may seem plausible for workers in early periods, it seems less so for workers entering non-employment in recent years.

benefits. This rise in invalidity benefits claimants appears to reflect workers with low employment prospects being systematically removed from the labour market.[49] The gap between long-term unemployment and invalidity benefits was largest from the mid 80s to 90s, when the Additional State Pension system provided supplementary payments. Invalidity benefits were simplified in the mid 90s, and the Additional State Pension system was abandoned. While there were no systematic efforts to provide incentives for workers receiving invalidity benefits to take up work again, labour market activation policies were increasingly adopted in the late 1990s, in reaction to the rise in inactivity.[50] Recall that changes in survival functions imply older workers entering non-employment in the 80s often never return to employment. Such large shifts to permanent joblessness disappear in later years. It stands to reason that the decline in medium skilled jobs, and the relatively generous provision of invalidity benefits to unemployment benefits, made it optimal for many workers, especially older ones, to leave the labour market and receive invalidity benefits until retirement. Shifts towards permanent joblessness for these groups entering in later periods may then have been reversed with policy changes starting in the mid 90s.

Apart from suggestive substantial explanations for the observed patterns, they may, of course, also reflect mismeasurement. For instance, instead of workers being permanently non-employed, observed spells could reflect workers returning from self-employment or emigration after long periods. While I cannot rule out these possibilities, it is suggestive to relate findings to aggregate changes in unemployment and inactivity population shares. Figures B.9 and B.10 show population share changes for employed, unemployed, and inactive workers derived from the Labour Force Survey for the same demographic groups.[51]

For male workers, we observe a sustained increase in inactivity for all age groups. The rise in inactivity peaks for older male workers in the mid 1990s, and decreases slightly thereafter. Younger and prime aged male workers experience increasing inactivity throughout the sample period. Unemployment is relatively high for all age groups in the 80s, but falls thereafter. There is no sustained increase in unemployment over the sample period.

Women experience a decline in inactivity over the sample period, although all age groups experience a temporary rise in the early to mid 80s. Young and prime

[49] Faggio and Nickell [2005] and Center of Economic Performance [2006] cite evidence that the Employment Serivce advised workers struggling to find employment to claim invalidity benefits, and that doctors were more likely to certify incapacity if the worker had a low probability of finding a job.

[50] See Faggio and Nickell [2005]; Center of Economic Performance [2006] for a discussion of the link between the rise in inactivity and policy changes.

[51] LFS data have been provided by Jennifer C. Smith. See Elsby et al. [2011].

aged women also experience temporary increases in unemployment, especially in the 80s, but no general increase in later years.

How do these patterns compare to changes in survival functions? On the one side, there is no direct link between changes in survival functions, based on flow sampling, and changes in annual population shares. Also, recall that results do not reflect workers entering non-employment during recessions. Nevertheless, the rise in inactivity for men seems largely consistent with generally longer non-employment spells. The fact that unemployment remains relatively constant may be taken to suggest that most of the longer spells, also the temporary ones, reflect inactivity rather than unemployment. The pattern over time is also roughly consistent with the observed time variation in shifts. Inactivity for older male workers peaks in the might 1990s, possibly reflecting the large inflow of workers remaining permanently jobless during the 80s, before policy changes occur. Young and prime age men see sustained shifts towards longer spells, which seems in line with the sustained increase in inactivity.

Results for women are also roughly in line with changes in population shares. Women of all age groups experience a tendency towards ever shorter spells. This seems consistent with the rise in the employment population share over the sample period. One may wonder, however, about relative magnitudes. After all, the push towards shorter spells is strongest for young women, whose employment share changes do not seem to differ drastically from other age groups. On the other side, the rise in the employment ratio starts later for older women, reflecting a prolonged rise in inactivity, lasting until the mid 80s. This would be compatible with the larger shifts towards permanent spells observed for older women during this period.

Overall, this chapter provides evidence which is suggestive for job polarization being associated with a general increase in non-employment spells, as the decline in demand for medium skilled jobs implies that workers take more time to reallocate to new jobs, and that some workers have become permanently jobless as a result. If these results hold true, this constitutes an important qualification of the narrative of Creative Destruction: while redundant jobs are replaced with new ones, some workers experience prolonged periods of joblessness as they are either unable or unwilling to secure available jobs. Variation across groups suggests that the failure to reallocate reflects both aggregate as well as worker-level factors. The fact that women experience much less adverse changes, and are even able improve their situation, while men seem to be hit much harder, begs the question which factors can explain the better performance of women. Understanding these factors is an important task for future research. This understanding is vital for improving the process of labour reallocation, and to avoid human suffering, in face of the possibly fundamental changes in labour markets still ahead of us.

Chapter 4

Does Skill-Biased Technological Change Differ Across OECD Countries?

4.1 Introduction

Technological change has important effects on the labour market. By affecting the productivity of workers, it determines their wage and employment prospects. Recent decades have witnessed substantial increases in wage inequality, which have been linked to SBTC.[1] SBTC affects labour productivity of skilled and unskilled workers differently, and so contributes to wage inequality by raising the skill premium. Additionally, SBTC has been linked to the differential rise in unemployment rates in the US and Europe in the 80s and 90s.[2]

Technological change, and SBTC in particular, have often been assumed to be exogenous shocks to the economy, implying that shocks are independent of country-specific factors. Given this independence, and the fact that OECD countries share similar technologies, it is further assumed that SBTC shocks are common to OECD countries.[3] This assumption of common SBTC has important implications on the choices policy makers face: if SBTC is an exogenous shock to the economy, policy makers can affect the impact on the labour market only within the boundaries imposed by these technological changes. The Krugman-hypothesis, for instance, posits that in face of SBTC policy makers can only choose between either higher unem-

[1]See Acemoglu and Autor [2011] for a review of the link between technological change and rising wage inequality.

[2]See Krugman [1994]; Blanchard [2006].

[3]See, for instance, Blanchard [2006]; Blanchard and Wolfers [1999].

ployment and or higher wage inequality.[4] If, however, SBTC is not exogenous to the economy and differs across countries, it stands to reason that these differences reflect country-specific factors which are, to some extent, under control of policy makers. Clearly, the scope of policy makers for influencing the impact of technological change on the labour market is larger in the latter case.[5] If, for instance, policy makers are able to guide innovation towards raising productivity of unskilled workers, the magnitude of SBTC, and the pressure on either rising unemployment or wage inequality, could be dampened.

Whether or not SBTC is constant across countries is therefore an important question. What is the basis for the assumption of common SBTC for OECD countries? As I argue in this chapter, it is not based on evidence. A large body of literature has established that rising wage inequality experienced in many OECD countries in recent decades in large parts does reflect SBTC. However, apart from the claim that SBTC is present in many OECD countries, little is known about comparative aspects of SBTC.[6] Most papers focus on US data and, based on the assumption of rapid technology diffusion, generalize implications about SBTC to other OECD countries.[7] Cross-country studies often assume constant SBTC from the start, and attribute any variation between countries to differences in relative skill supply changes or institutional differences, although they do not control for the latter. While the focus of these papers is to establish SBTC as an international phenomenon, they devote little attention to cross-country variation in magnitudes, which are frequently found.[8]

In addition to variation in demand-shift measures found in these studies, but not elaborated on, there are theoretical reasons for technology to vary across OECD countries: the literature on directed technical change and transaction costs based approaches stresses cross-country variation in technology factor mixes and institutional idiosyncrasies as potential determinants of cross-country variation in

[4]See Krugman [1994] for an exposition of this view.

[5]It is important to note that to date the literature falls short of establishing robust results for explaining cross-country unemployment patterns. See for an exposition of the 'shocks-and-institutions' hypothesis Blanchard [2006]; Blanchard and Wolfers [1999]. For assessments focusing on institutions, see Layard [1999]; Nickell et al. [2005]. For a critical review of the existing literature, see Baker et al. [2005].

[6]For more recent reviews for evidence on SBTC see Aghion et al. [2014] and Salverda and Checchi [2014].

[7]See, for example, Card and DiNardo [2002] Card and DiNardo [2002]; Katz and Murphy [1992]; Bound and Johnson [1992].

[8]See Caselli [2005]; Katz et al. [1995]; Berman et al. [1998]; Murphy et al. [1998]; Machin and Van Reenen [1998]; O'Mahony et al. [2007]; Gottschalk and Joyce [1998]. For a notable exception, see Card et al. [1999].

technology and technological change.[9]

The gap in the literature may be explained by the inherent difficulty in assessing SBTC, especially in a cross-country context. Despite improvements in recent years, data availability remains limited, and potential measurement errors are severe. Additionally, technology is defined in terms of knowledge, and is therefore not observed. As a result, technology is often identified in terms of residuals: unobserved demand shifts necessary to account for observed patterns in labour market outcomes are assumed to reflect technological changes. Of course, the validity of such conclusions is limited, as many factors varying across countries and over time are not controlled for.[10] Alternatively, technology proxies can be used, but good proxies of technology are hard to come by, and existing ones carry little comparative information for qualitative technological differences.[11]

Comparing SBTC across countries is therefore difficult, but the question whether SBTC is common or differs among OECD countries has important implications, and the case for common technology is not self-evident. Making assumptions informed by empirical studies, even with limitations, is epistemologically preferable to basing assumptions on ad hoc reasoning, especially when theoretical knowledge and empirical findings appear to contradict this reasoning. In this chapter I therefore ask whether SBTC varies across countries, and which factors might be associated with any SBTC differences found. In light of the difficulties for identifying SBTC mentioned above, I assume a rhetorical perspective: I extend the seminal empirical approach used to establish SBTC to a cross-country context, introduced by Katz and Murphy [1992], additionally controlling for confounding variables generally ignored in this literature. Using this approach, I ask whether results based on this approach are compatible with constant SBTC across a wide range of OECD countries. The implication is that, if one accepts conclusions on SBTC based on this approach, one ought to accept conclusions about the presence or absence of SBTC differences.

In particular, because technology is not directly observable, the empirical approach taken in this paper is to identify SBTC in terms of the partial correlation between wage inequality and relative skill supply. Graphically, this correlation

[9]See Acemoglu [2009]; Acemoglu and Zilibotti [2001]; Acemoglu [2003]; Fadinger and Mayr [2015]; Pischke [2005]; Huo [2015].

[10]See Katz and Murphy [1992] for one of the first expositions of this approach.

[11]For instance, Card et al. [1999] compare data on computer use for the USA, Canada and France. They argue in favour of the proxy but mention that computer use is substantially more widespread across less educated workers in France and Canada as compared to the USA. This could be taken to suggest that in France and Canada computer use raises productivity of less skilled workers relative to less skilled US workers. Clearly, using data on computer use as a proxy for skill-biased technological change would be misleading in this case.

indicates the slope describing the relationship between wage inequality and relative skill supply. That is, a positive partial correlation implies a positive slope, as higher relative skill supply is associated with higher relative wages. An increase in relative skill supply being associated with an increase, rather than a decrease, in wage inequality, suggests an outward shift in relative skill demand. In line with the literature, such demand shifts are assumed to reflect SBTC. This reasoning suggests that higher SBTC implies a steeper positive slope between relative skill supply and relative wages. Controlling for confounding variables, and assuming that demand shifts reflect SBTC, observing a steeper slope in country A than in country B indicates that technological change in country A is more skill biased than in country B. An important determinant of wage inequality, which is often not controlled, for is institutional change. The literature on Varieties of Capitalism (VoC) stresses the importance of institutional complementarities, suggesting that effects of institutions are difficult to identify in isolation.[12] I therefore use factor variables, derived from various institutional measures, to control for the institutional environment, including interaction terms to allow for institutional interactions.

Finding significant SBTC differences across countries, I further extend this approach and ask which country-specific factors might explain this variation. As institutions have been suggested as possible explanations, I examine whether SBTC differences are systematically related to the institutional measures described above. Results provide some tentative evidence that SBTC differences are indeed related to institutional measures, but other country-specific factors, not identified by time-varying institutional measures, appear to be important as well. Findings therefore suggest substantial scope for policy makers to affect labour market outcomes even in presence of SBTC, although examining particular channels is beyond the scope of this chapter.

My contributions to the literature are as follows: this chapter is the first in using the approach suggested by Katz and Murphy [1992] to focus on the assumption of common SBTC. It is also the first to examine in such a context whether SBTC differences are associated with institutions. It thereby contributes to the above cited literature on cross-country determinants of wage inequality and unemployment by testing, and rejecting, a widely used assumption. Moreover, it contributes to this literature by exhibiting links between technological and institutional explanations. Findings are suggestive for future research to further explore the scope policy makers have to affect the impact of technological change on the labour market. Findings are also informative for the literature on VoC, which has focused on technology

[12]See Hall and Soskice [2001].

differences in terms of incremental versus radical innovation, or product versus process innovation. The implications of this chapter are that technology differences related to the institutional environment may extend to the skill bias. Finally, it contributes to the literature on directed technological change by examining patterns of yet undocumented technological variation and by providing evidence for institutional determinants of technological change.

The remainder of the chapter is as follows. Section 4.2 discusses the seminal empirical contributions to SBTC, focusing on the implications of this literature on SBTC variation. Also, I briefly discuss theoretical arguments for SBTC differences and potential determinants of this variation. Section 4.3 discusses the empirical analysis and results. Section 4.4 concludes.

4.2 The Consensus View on Skill-Biased Technological Change

The aim of this chapter is to assess whether seminal methods used to establish SBTC are consistent with the assumption of common SBTC. In this section, I therefore discuss the seminal literature on skill-biased technological change to motivate the empirical approach taken in the rest of the chapter. First, I discuss the concept of SBTC, and how it has been investigated empirically. Second, I critically assess the evidence with the view on SBTC differences across countries. Third, I review some theoretical arguments for country-level variation in technology, with particular focus on variation in SBTC.

4.2.1 Identifying Skill-Biased Technological Change

In this subsection I define the concept of SBTC. This concept is defined in terms of its effects as follows: in the context of an aggregate production function with constant elasticity of substitution, technological change is skill-biased if the marginal effect of an increase in technology on the relative marginal product of skilled to unskilled workers is nonnegative. More formally, following Acemoglu [2009], for a production function $Y = F(H, L, A)$, technological change is skill-biased if

$$\frac{\partial \left(\frac{\partial F(H,L,A)/\partial H}{\partial F(H,L,A)/\partial L} \right)}{\partial \left(A \right)} \geq 0$$

with A denoting technology, Y output, H skilled and L unskilled labour. The definition of skill is open to interpretation and varies in the literature, but most commonly the distinction between skilled and unskilled workers is drawn by college attendance, whereby a worker qualifies as skilled if she has at least some college education or higher degrees, and as unskilled otherwise.[13] Importantly, skills can refer to either observable or unobservable characteristics. However, while there is substantial unexplained variation within observable groups, I follow the bulk of the literature and address observed differences in educational attainment only.

Technology is commonly defined as knowledge, i.e. the knowledge of how to combine and organize inputs to produce something desirable, and so is generally not observed directly.[14] As result, the literature attempts to identify SBTC in terms of its effects. Any empirical investigation therefore faces the problem of relating observed changes in the relative marginal product of skilled and unskilled workers to unobserved changes in technology.

Several strategies have been suggested. In the literature on aggregate outcomes, three different approaches can be distinguished. All of them start from a simple demand and supply framework with workers of different skill groups being imperfect substitutes. Consequentially, changes in relative labour market outcomes, such as relative wages or employment or the skill premium, can be thought of as the result of changes in supply or demand curves for different worker groups. SBTC affects relative labour market outcomes via relative skill demand. I discuss the three approaches in turn.

Katz and Murphy Approach

The most widely used approach identifies SBTC with residual changes in relative labour market outcomes after controlling for relative skill supply.[15] More formally, assume the aggregate production function takes the following form:

$$Y = \left[\alpha \left(A_H H\right)^{\frac{\sigma-1}{\sigma}} + (1-\alpha)\left(A_L L\right)^{\frac{\sigma-1}{\sigma}}\right]^{\frac{\sigma}{\sigma-1}}$$

where Y is aggregate output, H and L are skilled and unskilled workers, and A_H and A_L denote technologies complementing skilled or unskilled workers respectively, with elasticity of substitution between skilled and unskilled labour being $\sigma > 1$. Assuming a perfectly competitive labour market we can define the skill premium ω as the ratio of marginal products of skilled to unskilled workers as

[13] Alternative measures often used are based on distinguishing non-production versus production workers, whereby the former refer to skilled and the latter to unskilled workers, or by distinguishing workers by occupation based on their hypothesized skill requirements.

[14] See Mokyr [1990].

[15] For the very first exposition of this approach see Katz and Murphy [1992].

$$\omega = \left(\frac{1-\alpha}{\alpha}\right)\left(\frac{H}{L}\right)^{-\frac{1}{\sigma}}\left(\frac{A_H}{A_L}\right)^{\frac{\sigma-1}{\sigma}}$$

If $\sigma > 1$, that is if skilled and unskilled workers are imperfect substitutes, an increase in technology, as defined by an increase in A_H/A_L, raises the skill premium and is therefore said to be skill biased. In logs, the expression becomes

$$\log\omega = \log\left(\frac{1-\alpha}{\alpha}\right) - \frac{1}{\sigma}\log\left(\frac{H}{L}\right) + \frac{(\sigma-1)}{\sigma}\log\left(\frac{A_H}{A_L}\right) \tag{4.1}$$

Based on equation 4.1 one can identify the skill-bias in two ways. First, solving this expression for $\log(A_H/A_L)$ and using data on the skill premium, relative skill supply, and the structural parameters of the production function, one can calculate the skill bias of technology. To get SBTC, one can simply compute the above expression in first differences. Second, as done in the original paper, assuming smooth SBTC, one can regress the skill premium on a constant, relative skill supply, and a time trend, retrieving the average skill bias over the period from the estimated time trend coefficient. The regression equation corresponding to the above expression is:

$$\log\omega_t = \beta_0 + \gamma t + \beta_2\log\left(\frac{H}{L}\right)_t + \varepsilon_t \tag{4.2}$$

As this approach does not control for other inputs, demand shifts identified by the residual reflect changes in prices and supplies of any inputs not explicitly accounted for in the above production function. Either approach can be said to identify SBTC in residual terms, as inference about SBTC is based on interpreting the discrepancy of observed changes in relative skill supply in accounting for observed changes in skill premia as indicating SBTC.

The validity of this approach is based on several assumptions. The confidence with which one can identify relative demand shifts with SBTC rests on controlling for all confounding variables. Several explanations are possible for the observed changes in relative wages in a simple demand and supply framework: changes in relative skill supply, changes in relative demand, and, by extension, institutional changes directly affecting factor prices. Demand changes in turn can be due to several factors, the most important mentioned in the literature being trade and SBTC. Note that although the combined evidence suggests that trade only plays a minor role, most studies also cite institutions as important alternative explanations, and several studies examining the role of institutional change on wage inequality, even for the US alone, have found these to be important.[16] This suggests that studies

[16]See Bound and Johnson [1992]; Freeman and Katz [1995]; Katz et al. [1995]; DiNardo et al.

examining SBTC based on such a supply and demand framework should control for relative skill supply, trade intensification, and institutions. If any of these variables are omitted, identifying SBTC via unexplained changes in the dependent variable does not actually identify SBTC, but some combination of SBTC and the omitted variable.

Shift-Share Analysis

Alternatively, one can use a shift-share analysis to examine the sources for changes in relative skill demand. Specifically, using industry-level data, one can decompose aggregate changes in relative labour market outcomes, such as employment or wage bill shares, into changes within and between industries. It is then hypothesized that changes in favour of more skilled workers reflect demand shifts due to SBTC if they take place within industries, whereas such changes reflect demand shifts due to, say, trade intensification, if they take place between industries.[17]

Note that shift-share analyses relate within-industry shifts to SBTC without controlling for changes in relative skill supply. It is, of course, assumed that relative skill supply is increasing at the same time. The reasoning is that if the aggregate wage bill share for high skilled workers increases because all industries use relatively more skilled labour, despite increasing relative skill supply, one may conclude that factors raising relative demand for skilled labour are sufficiently general to apply to all industries. Technology, in contrast to trade, is thought to exhibit such a general effect across industries. Changes reflecting between-industry shifts are assigned to competing explanations. A shift towards industries using relative more high skilled labour, for instance, is presumably driven by trade. While such decompositions are suggestive about driving factors, their inability to control explicitly for relative skill supply conceptually limits their ability to assess and compare the magnitude of SBTC across countries. Also, not all competing explanations can be assigned to between-industry shifts. For instance, if institutional change affects workers with high and low wages differently, e.g. via deunionization, within-industry shifts may partly reflect institutional change.

Technology Proxies

Finally, one can use proxies for technology, e.g. the fraction of workers using computers, research and development (R&D) expenditures, or ICT capital intensity, to examine demand shifts more directly. Three approaches have been used in the

[1996]; Card and DiNardo [2002].

[17] For one of the first expositions of this approach see Bound and Johnson [1992].

literature. First, one can use augmented wage equations to assess the impact of ICT usage at the micro-level. Second, starting from cost functions, one can derive expressions relating labour market outcomes to observable inputs and unobserved labour demand, and based on the hypothesis that unobserved demand changes are driven by technological change, use proxies for technological change to control for these unobserved demand changes. Third, one can test whether these proxies take up substantial parts of the variation in changes in relative labour market outcomes.[18]

The use of such proxies for technology differences across countries is limited, however, if proxies and particular technologies are not directly linked. Recall that technology is defined as knowledge. Technology can therefore be associated with, but not reduced to, ICT capital goods. For instance, the technology embodied in ICT capital refers to the knowledge how to produce and to employ this type of capital in the wider production process. Using the same ICT capital goods differently amounts to using a different technology. The link between measures for ICT capital and technology may therefore differ across countries and change over time as ICT capital goods are used in different ways.[19] In this context, measures for ICT capital are not equivalent to SBTC measures across countries because the use of ICT capital by skilled and unskilled workers can differ across countries.[20]

4.2.2 Empirical Findings on Skill Biased Technological Change

Using the approaches described above, the combined findings of studies establish the consensus view of SBTC being the main driver for increasing wage inequality for between-group wage inequality in OECD countries during the last decades. In this section, I discuss the findings giving rise to this view more closely.[21]

[18]For an approach using micro-level evidence attempting to assess causal effects see Krueger [1993]. For an approach using macro-level evidence focusing on correlations, see Huo [2015].

[19]See Griliches [1998] for a similar argument.

[20]As mentioned above, Card et al. [1999] find that computer use is substantially more widespread among unskilled workers in France and Canada than in the USA.

[21]For an early discussion of the consensus view focusing on US evidence see Card and DiNardo [2002]. Also, note that the consensus view of SBTC driving wage inequality has been superseded by the more encompassing hypothesis of routine-biased technological change, of which SBTC can be seen as a special case. On this account, the rise in wage inequality in past decades is captured well by the dichotomy of skilled and unskilled workers underlying SBTC, whereas employment changes, especially in more recent years, are better captured by distinguishing between more skill types. See, for instance, Autor et al. [189-194]; Goos et al. [2014]. I focus on SBTC for reasons of data availability, while keeping in mind that findings may not extend beyond the period covered by the sample.

US Evidence

For the US, based on the approach advocated by Katz and Murphy [1992] or shift-share analyses, a large number of studies have found that increases in relative skill demand are necessary to explain increasing wage inequality between educational groups, occurring in spite of simultaneously increasing relative skill supply, that overall shifts are mostly due to within-industry variation, and that SBTC would be the remaining major explanation compatible with these within-industry shifts and rising skill demand.[22] Approaches using proxies confirm that workers using computers have higher wages, and that the pattern of computer use can explain a substantial part of the increase in educational wage gaps (Krueger [1993]), and that there is a statistically significant relationship between increased computer use and faster skill upgrading at the industry-level, which appears to have increased over time (Autor et al. [1998]; Berman et al. [1994]).

International Evidence

Several studies have sought to generalize the findings for US data to other OECD countries. Some of these widely cited studies are based on straightforward manual comparisons of changes in relative skill supply and relative wages (e.g. Katz et al. [1995] and Gottschalk and Joyce [1998]), or decompositions of overall changes into within- and between-industry shifts (Berman et al. [1994], Machin and Van Reenen [1998]). The observed patterns are largely consistent with demand shifts, mostly due to within-industry variation, being present in most countries. However, Gottschalk and Joyce [1998] find only weak evidence for educational wage gaps (as opposed to experience wage gaps), with contradicting findings for some countries. Berman and Machin [2000] extent this analysis to middle- and low-income countries and find that most countries show patterns consistent with SBTC since the 1980s. Murphy et al. [1998] and Acemoglu [2003] explicitly replicate the approach based on Katz and Murphy [1992] for a larger number of countries. Findings suggest that constant SBTC across countries can explain patterns in educational wage gaps well for some but not for all countries. Studies relying on proxies find evidence largely consistent with SBTC being an international phenomenon for many OECD countries, although findings leave much room for considerable variation across countries (Machin and Van Reenen [1998]; Card et al. [1999]; O'Mahony et al. [2007]).

Note that these papers all focus on SBTC. For a purely cross-sectional analysis focusing on the level of skill-bias, Caselli and Coleman II [2006] follow the ap-

[22]Seminal papers are Bound and Johnson [1992]; Levy and Murnane [1992]; Berman et al. [1994]; Freeman and Katz [1995]; Katz and Autor [1999].

proach of Katz and Murphy [1992] using cross-sectional data from the 1980s to back out technologies complementing skilled or unskilled labour. Their results (p. 515, figure 5) show OECD countries producing at or close to the technological frontier, but exhibiting a large dispersion in terms of relative efficiencies. For instance, the US is among the countries using skilled labour most efficiently and unskilled labour most inefficiently, whereby Germany, France and Italy are using skilled labour just as efficiently, but are among the countries using unskilled labour most efficiently.

Limitations

These findings have to be seen in light of the studies' limitations. As discussed above, unless studies control for all confounding variables, identifying SBTC with unobserved demand shifts necessary to explain observed changes in relative labour market outcomes may confuse SBTC with any of the omitted variables. Confounding variables mentioned in the literature are relative skill supply, institutions and trade openness. However, most of the studies cited above do not control for either institutions, trade, or both.

In particular, except for Bound and Johnson [1992] and, to some extent Card et al. [1999], none of the above studies control for institutions. Adopting a cost-function approach, Berman et al. [1994], Machin and Van Reenen [1998] and O'Mahony et al. [2007] also do not control for relative skill supply, although Machin and Van Reenen [1998] explicitly say they should. In addition, many studies do not control for trade intensification (Autor et al. [189-194]; Gottschalk and Joyce [1998]; Berman et al. [1998]; Murphy et al. [1998]; Machin and Van Reenen [1998]; Card et al. [1999]; Berman and Machin [2000]; Acemoglu [2003]; O'Mahony et al. [2007]).

Bound and Johnson [1992] attempt to assess all alternative explanations in a unified framework. However, they identify trade effects as between-industry changes and, due to a lack of data, they identify institutional change as sector-specific wage effects common to all worker groups in this sector. Card et al. [1999] attempt to examine the role of institutions by testing parameter restrictions derived from their model. For each country, they regress changes in wages and employment across different groups on labour supply shocks and proxies for labour demand, arguing the Krugman-hypothesis (Krugman [1994]) suggests the effect of technology proxies should vary with countries' institutional environments. However, they do not explicitly control for institutions or institutional change. While they find heterogenous effects across countries, their findings are largely inconsistent with the Krugman-hypothesis. Caselli and Coleman II [2006] mention institutions but argue that, in the context of development accounting where rich countries would tend to have more egalitarian labour market institutions, omitting institutions would bias results against their finding.

Reassessing Common SBTC

What can we conclude about SBTC based on these studies? The combined evidence from studies using the different approaches mentioned above suggests that SBTC played an important role for increasing wage inequality in many countries. This wide acceptance of SBTC as a major determinant of wage inequality increases rests on the large increases in wage inequality occurring in many countries and the robust findings using various approaches.

Beyond this claim, however, the evidence is much weaker. In fact, most of the studies above address the question whether SBTC has been driving wage inequality increases, or whether SBTC is an international phenomenon. As such, they try to answer the question whether SBTC has been present in the US or other OECD countries. However, they do not address the size of the effect of SBTC on wage inequality, absolutely or compared to other factors such as trade or institutional change. Neither do they address whether the size of SBTC varies across countries.

Studies using the approach by Katz and Murphy [1992] need to control for all confounding variables to answer questions about the size of the effect or its relative importance. Because none of them control for institutions, although institutions have been found to affect wage inequality, and some not even for trade openness, they are unable to provide evidence about the size of SBTC. Comparing SBTC across countries requires controlling for confounding factors varying across countries: because both institutions and trade openness vary across countries, these studies cannot tell whether SBTC is similar across countries. Of course, studies control for some confounding variables at a time. Some studies control for trade and find its effect to be quantitatively small. Others argue that institutional change may be important but too small to explain the changes in wage inequality. However, because SBTC is essentially identified as residual, and because findings are biased when omitting relevant variables, unless all competing explanations for rising wage inequality are assessed in a unified framework one cannot tell whether residuals leave room for SBTC. This is particularly important when assessing SBTC internationally, as the relative importance of institutional change varies substantially across countries.

An additional limitation is that most of the studies cited above examining SBTC as an international phenomenon only cover a small range of countries. For instance, Katz et al. [1995] only focus on the US, UK, Japan and to lesser extent on France. Murphy et al. [1998] only look at the US and Canada, Card et al. [1999] only on the US, Canada and France, Greiner et al. [2004] only compare the US and Germany and O'Mahony et al. [2007] focus on the US, UK and France. Larger samples were used by Gottschalk and Joyce [1998], Berman et al. [1998], Machin and Van Reenen [1998] and Acemoglu [2003]. Yet, neither of these studies even attempted to control for institutions.

Studies using the shift-share approach can provide suggestive evidence for the presence of SBTC in a particular country, to the extent it is associated with within-industry shifts. The lack of explicit control for relative skill supply or competing explanations, however, means this approach does not allow quantifying SBTC and is therefore unable to assess the relative importance of SBTC across countries. Similarly, as discussed above, the proxies used for technological change provide only a weak link to particular technologies, i.e. technology's skill bias, and thus are ill-suited to compare SBTC across countries.

In sum, while we know that SBTC occurs in many OECD countries, we know much less about the magnitude of SBTC, or about comparative aspects of SBTC. Consequentially, the seminal literature does not provide evidence in favour of the hypothesis that SBTC is constant across countries. Instead, studies often assume from the start that SBTC is constant across countries (see Freeman and Katz [1995]; Gottschalk and Joyce [1998]; Berman et al. [1998]; Murphy et al. [1998]; Katz and Autor [1999]; Acemoglu [2003], although the latter clearly uses this assumption rethorically). This assumption is made despite the fact that studies examining cross-country data consistently show substantial variation in magnitude, and sometimes even in sign, either in terms of within- and between-industry variation (Katz et al. [1995]; Berman et al. [1998]; Berman and Machin [2000]), cost-function estimates using proxies (Machin and Van Reenen [1998]; Card et al. [1999]), estimates of elasticity of substitution (Greiner et al. [2004]), or residual estimates of SBTC (Acemoglu [2003]).

If one wants to examine common SBTC as an international phenomenon of similar magnitude, it is preferable not to use shift-share analyses or rely on proxies. Instead, if one wants to identify SBTC in residual terms, one should use a unified framework to control for all competing explanations. Furthermore, one should use samples covering a range of diverse countries.

Of course, an empirical analysis facing these difficulties may carry less weight if there are strong theoretical reasons to believe that SBTC in one country extends to similar SBTC in others. In the next section I address this concern by reviewing some theoretical arguments supporting the notion that technology, and SBTC in particular, does vary across countries.

4.2.3 Theoretical Arguments for Country-level SBTC Variation

Much of the literature discussed above implicitly or explicitly assumes that technology, or technological change, is constant across countries. In this section, I briefly review some arguments why SBTC may vary across countries. As I will discuss below, institutions varying at the country-level are suggested as determinants of SBTC. However, recall that the aim of this chapter is not to identify particular institutional mechanisms. Instead, I want to test whether SBTC differs across countries, and whether this variation is related to institutional measures. Accordingly, this review focuses on arguments that technology may differ across countries, and that it may do so systematically with institutions, rather than providing a comprehensive review how exactly institutions affect technology.[23] I consider three strands of literature, partly overlapping, which provide reasons for SBTC to vary at the country-level: directed technical change, labour market imperfections, and Varieties of Capitalism.

First, consider models on directed technical change. The baseline model examines the allocation of R&D effort between technologies complementing either skilled or unskilled labour (Acemoglu [2009]). R&D firms can choose to either enter the market for technologies complementing skilled or unksilled labour, both governed by a free-entry condition. The stock of technology complementing a particular labour type depends on the number of firms entering this market. The more profitable it is to make an innovation, the more firms enter and innovate. The ratio of skilled and unskilled labour complementing technologies, the skill bias, is thus determined by the relative profitability of innovating in either market. This profitability in turn is determined by a price and a market effect: the price effect reflects the fact that innovations complementing a factor input will be more profitable if this factor is scarce. The market effect reflects the fact that innovations complementing a factor input will be more profitable if the market for this factor is larger, i.e. the

[23]Note that reasons for constant technology, as opposed to SBTC, are implied by the literature on development accounting. This literature typically conceptualizes technology as total factor productivity (TFP), and aims to examine why TFP differs across countries. Assuming a common technology frontier defined by the currently most efficient technology, the distance to the frontier is determined by barriers to technology adoption. Technology differences result from differences in institutional quality: the more institutions impinge on property rights, the weaker are incentives for adopting new technologies (Caselli [2005]; Hall and Jones [1999]; Jerzmanowski [2007]). The implication is that OECD countries, with similar institutions, are all close to the frontier. However, this only applies to technology understood as TFP. The literature on appropriate technologies suggests that countries may differ in the relative efficiency with which they use inputs, while possibly being equally efficient overall (Acemoglu and Zilibotti [2001]; Jerzmanowski [2007]).

factor is more abundant. Overall, an increase in relative skill supply raises the skill bias if the elasticity of substitution between skilled and unskilled labour is sufficiently large.[24] Directed technical change therefore suggests two determinants for the skill bias: relative skill supply and relative factor prices. It stands to reasons that, on this account, other determinants affecting the relative price or employment of skilled labour also determine the skill bias. Candidate factors are trade openness and institutions. The channel introduced by directed technical change for trade acts to exacerbate the effect already discussed: a country abundant in high skilled labour opening to trade with countries abundant in low skilled labour shifts production to high skilled labour, making high skilled labour more scarce. An example for the effect of labour market institutions is discussed next.

Fadinger and Mayr [2015] introduce labour market institutions into a model of directed technical change. They combine a static model of directed technical change with technology adoption with a search and matching model in the labour market for monopolistically competitive R&D firms. The model features two types of frictions: search frictions are introduced as hiring costs, and institutional frictions are introduced as unemployment benefits and firing costs. For their baseline calibration they find a non-monotonic relationship between relative skill supply and the skill bias, whose turning point depends on the level of institutional rigidities: institutional rigidities introduce a wage floor for low skilled workers. If this wage floor is not binding, an increase in relative skill supply increases the skill bias as in the baseline model for directed technical change. If the skill bias is sufficiently large, wages for low skilled workers hit the wage floor, below which wages for low skilled workers cannot fall. Further increases in relative skill supply act to diminish the skill bias. For a given relative skill supply, the skill bias is decreasing with the level of institutional rigidities.[25]

Second, other models have focused on labour market institutions in the con-

[24]In the baseline model an increase in relative skill supply is skill biased if the elasticity of substitution between skilled and unskilled is greater than two. The market effect then dominates the price effect.

[25]The literature on directed technical change identifies market forces and price setting institutions as determinants for qualitative differences in technology. However, some limitation of this literature are worth pointing out. First, because it assumes only two different labour types, it is not able to explain job polarization. Second, the baseline dynamic model of directed technical change implies an increasing skill premium but, in absolute terms, wages for unskilled workers increase. This is in contradiction with US evidence. Third, in order to explain skill-biased technological change by increasing relative skill supply, the model relies on an elasticity of substitution which is at the border of available estimates. A consensus estimate is between 1.4 to two. The baseline model requires a value greater than two. The model by Fadinger and Mayr [2015], which adds technology adoption and introduces a non-vertical labour supply curve reflecting search costs, requires an even higher estimate.

text of frictions giving rise to rents. Labour market institutions can induce firms to raise productivity of low skilled workers if institutions introduce a binding wage floor for low skilled workers. Using a search and matching framework, Acemoglu [2003] provides a model where firms observe the worker's human capital type as being either low or high after being matched, and subsequently choose whether or not to adopt technology at some fixed cost, whereby the productivity gain associated with technology adoption is increasing in the worker's human capital. Bargaining over the match surplus takes place after technology adoption costs are sunk, and workers and firms retain fixed shares. If the cost of adopting technology is sufficiently low, firms matched with either worker maximize profits by adopting technology. There exists a threshold for adoption costs after which firms choose to adopt technology only if they are matched with high human capital workers. The skill premium is related to the adoption decision: higher adoption costs, implying that firms matched with low skilled workers choose not to adopt, imply a higher skill premium. Institutions are introduced in terms of a minimum wage, which is binding only for firms matched with low human capital workers if they choose not to adopt technology. Depending on the level of minimum wages, firms adopt technology for low skilled workers with but not without minimum wage. Firms matched with high human capital workers always adopt technology. Sufficiently high minimum wages can thus decrease the skill-bias. Intuitively, the gain from adopting technology for low skilled workers increases with minimum wages as firms already pay high wages, and can claim a larger share of the increase in productivity.

This argument is further elaborated on in the literature on training in Pischke [2005]. Similar to Acemoglu [2003], Pischke [2005] argues that in the presence of labour market frictions leading to wage bargaining, institutional rigidities can induce firms to invest in training of unskilled workers. In his graphical example, firms enjoy rents as frictions imply that wages are below the worker's productivity. Productivity $f(t)$ and wages $w(t)$ are increasing in skills, and skills are equivalent to training t. Rents for employing a worker with skill-level t are $\pi(t) = f(t) - w(t)$. Workers differ in their skill levels, and firms can invest in training and raise the worker's productivity at some fixed costs. However, rents are independent of the level of training, as training is assumed to raise productivity and wage in equal amounts, implying the firm has no incentive in paying for training of any worker. Minimum wages w_{min} alter the firm's incentives: the minimum wage is binding for workers with low skill levels. Firms employing such workers can raise profits if the increase in productivity exceeds training costs. The gain in profits is larger the larger the discrepancy between the minimum wage and the worker's initial skill level. Again, this implies firms invest more if the minimum wage is higher. Pischke also suggests to interpret training as investments in physical capital embodying technology. In this case, firms invest more in training or technology, raising the productivity of unskilled workers, the more rigid institutions are.

Third, the literature on varieties of capitalism discusses more broadly the effect of institutions and institutional complementarities at the country level on economic activity, including innovation. The approach is based on the literature viewing transaction costs as ubiquitous, arguing that attempts to economize on these transaction costs lead to organizational structures and institutions of various sorts, starting from the existence of firms, extending to relationship between firms, and governmental institutions (see Coase [1937]; Williamson [1985]). VoC takes as a conceptual starting point the firm, which is seen as the key agent for adjustment in face of technological change or trade intensification.[26] In order to produce, firms interact with various agents, both internally with its own employees and externally, e.g. with other firms or stakeholders. Because of asymmetric information, implying problems of moral hazard, adverse selection, and shirking, these relationships are characterized by strategic interactions. Institutions can affect these strategic interactions, for instance by providing information about other agents' preferences, or expected payoffs, or by affecting their structure, allowing it to monitor other agents' actions. Institutional setups vary at the country level, so firms may find it optimal to pursue different strategies in different countries. In particular, Hall and Soskice [2001] suggest two ideal-types, which may be thought of as end points on a continuous scale: liberal market (LMEs) and coordinated market economies (CMEs). LMEs organize strategic interaction predominantly via hierarchies within firms, and relying on competitive market arrangements, based on formal contracting, for relations between firms. Coordination between firms and in the labour market is primarily guided by price signals. For CMEs, on the other hand, coordination among firms, as well as between firms and their employees, often results from non-market interactions rather than price signals. Relationships are collaborative rather than competitive. Collaboration implies that agents coordinate on equilibrium strategies which offer higher returns to everyone involved. These strategies are available to firms in CMEs, more than to firms in LMEs, because of the presence of institutions facilitating the sharing of information, deliberation, as well as monitoring and sanctioning to prevent defection, e.g. widespread membership in business and industry associations relevant for standard setting and skill provision, wage setting and collaborative research. Importantly, Hall and Soskice [2001] stress the importance of institutional

[26] See Hall and Soskice [2001] for a seminal introduction to VoC.

complementarities.[27] Institutional complementarities are important because they imply that countries gravitate towards either ideal-type, rather than converging to a single institutional setup, so that institutional variation is persistent. Countries do not converge to the same institutional setup because, given institutional complementarities, the presence of some prior institutions affects the returns of other institutions. The initial institutional setup, itself based on historical contingencies, therefore provides agents in a particular country with reasons to further specialize the institutional setup, and firms to choose strategies supported by the institutional environment.

This also affects innovation. For instance, Hall and Soskice [2001] argue that CMEs possess a comparative advantage in producing differentiated goods based on incremental innovation. Such corporate strategies require a specialized workforce with firm-specific skills. To induce workers to acquire these skills, workers must perceive this investment as profitable. CMEs possess several, complementary institutional features supporting such investments. Firm-specific skills yield higher payoffs if workers enjoy long tenures, and if their current employer continues to pay appropriate wages. In CMEs, relatively high levels of employment protection assure workers that they are likely to enjoy long tenures, and wage negotiation is conducted by trade unions and business associations at the industry-level, so that wages for a given skill level are similar across employees. While such strategies are important for acquiring the skilled workforce necessary for incremental innovation, they require firms to be able to retain a stable workforce during economic downturns. Financial markets in CMEs complement these strategies, as membership in business associations, cross-ownership, and long-term relationships between firms and banks allow access to finance based on knowledge about the firm's longer-term profitability, rather than their current economic performance. The institutional setup in LMEs, on the other hand, is more conducive to radical than incremental innovation. Firms generally face weaker employment protection, set wages with individual workers rather than by collective bargaining, and access finance based on publicly available information about their current profitability. Instead, firms in LMEs can adapt more quickly to changes in the economic environment. For instance, technology transfer is facilitated more often by mobile labour. Managers exercise considerable authority over the corporate strategy, and thus are able to change strategies or adjust the

[27] They define two institutions as complementarity 'if the presence (or efficiency) of one increases the returns from (or efficiency of) the other' (Hall and Soskice [2001], p. 35).

workforce easily.[28][29]

To summarize, the models by Acemoglu [2003, 2009], Pischke [2005], and Fadinger and Mayr [2015] suggest that labour market institutions, by introducing sufficiently high wage floors, induce firms to adopt or innovate technologies raising the productivity of unskilled workers relative to skilled workers. Pischke [2005] and Fadinger and Mayr [2015] explicitly argue that the relative productivity of unskilled to skilled workers is increasing in the degree of labour market rigidities. Hall and Soskice [2001] and Huo [2015] provide arguments how the institutional environment can affect the type of technological change in general.

The literature on SBTC therefore identifies several potential determinants of SBTC at the country level: relative skill supply, i.e. the relative market size and scarcity of skilled and unskilled workers, and institutions, affecting relative prices of skilled and unskilled labour. The literature on VoC argues more generally for technological variation across countries. Moreover, the literature on VoC stresses that the effect of a particular institution, because of complementarities, depends on other institutions. Institutional interactions are important.

In the empirical section, I examine whether SBTC varies along these variables. As I argued above for using a unified framework, this suggests controlling for relative skill supply, institutions, and trade openness. To account for institutional complementarities, I measure the institutional environment in terms of factor variables, based on a range of institutions, and include interaction terms for these institutional measures.

4.3 Testing For SBTC Differences

I argued in the previous sections that comparative SBTC has not been addressed appropriately in the literature. The question if, and how, SBTC varies across OECD

[28] For another example, see also Huo [2015], who argues that more coordinated economies tend to engage more in process innovation, whereas more liberal market economies engage more in product innovation.

[29] The distinction between incremental and radical innovation does not relate directly to SBTC. However, the literature on VoC is suggestive for differences in SBTC. In CMEs, firms' actions, including technology adoption or innovation, depend more on collaboration than price signals, implying some scope for firms to raise productivity of low skilled workers to assure their collaboration. Suppose, for instance, technology adoption implies a strategic interaction between low and high skilled workers as well as the firm management, similar to a 'battle of sexes' game: Everybody benefits (assuming technology complements both worker types, albeit to varying degrees), but gains are differently distributed. Unskilled workers get relatively high wages, so - in line with the above logic - firms have incentives to invest in unskilled workers as they get a larger share of the productivity gain.

countries remains largely unanswered. In this section, I aim to address this gap. In particular, I examine two questions: first, extending the approach by Katz and Murphy [1992], I ask whether SBTC is constant across a large group of OECD countries. Second, I ask whether any differences in SBTC among these countries are related to labour market institutions.

Recall that technology is not directly observable. Empirical approaches typically identify SBTC in terms of its effects on some labour market outcome after controlling for confounding factors. In particular, in line with Katz and Murphy [1992], I identify SBTC in terms of the partial correlation between log relative earnings and log relative skill supply, which is equivalent to the elasticity of relative earnings with respect to relative skill supply. Graphically, this correlation is described by a slope. The simple demand and supply framework employed by Katz and Murphy [1992] suggests that an increase in relative skill supply decreases relative wages for a constant labour demand curve. An outward shift of the relative labour supply curve accompanied by an increase in relative wages implies an outward shift of the relative demand schedule. The larger the outward shift, the larger the increase in relative wages for a given increase in relative skill supply. If one identifies shifts in the relative demand schedule with SBTC, after controlling for competing explanations, an upward sloping relationship between relative wages and relative skill supply implies SBTC. A steeper relationship, implying larger outward shifts of the demand schedule, reflects larger SBTC. A positive partial correlation between measures for relative skill supply and relative wages can therefore serve as a measure for the magnitude of SBTC: observing a larger positive partial correlation, i.e. steeper slope, in country A than in country B, after controlling for all confounding variables, indicates that technological change in country A is more skill biased than in country B.

To exemplify the relationship examined in the remainder of the chapter, figure 4.1 plots the log of relative earnings against the log of relative skill supply, both in terms of deviations from country-specific means, for the countries contained in the sample. The line of fitted values, indicating the slope, is shown as a black solid line. Two features of the data can be observed: first, all countries in the sample experienced increasing relative skill supply over the sample period. Second, although most countries exhibit an upward-sloping relationship, it generally differs in magnitude and is negative for some countries.

It remains to be seen whether the differences shown in figure 4.1 are statistically significant, or whether they reflect changes in country-specific factors other than SBTC, and whether any remaining differences may be related to institutional measures. I do so in the following analysis. First, I examine whether country-specific deviations from the average slope are statistically significant after controlling for competing explanatory variables for changes in wage inequality. Second, I examine whether the slope varies significantly with institutional measures. The discussion in previous sections suggests controlling for institutions and trade openness. In particular, I regress the log of the earnings ratio, a measure for the skill premium, on the log of relative skill supply, institutional measures, a measure for trade openness, country-fixed effects and some additional control variables. In the main specification, I include the employment to population ratio to control for differences in unemployment and labour force participation. In additional specifications conducted as robustness tests I also control for relative educational expenses. To increase efficiency of the estimators in face of collinearity, and to capture the effect of institutions exhibiting complementarities, I use factor analysis to summarize the information contained in institutional variables into factor variables, which are then used as measures for the institutional environment.

Before discussing results for the empirical analysis, I briefly describe the data used for the main analysis.

4.3.1 Data

In this section I briefly describe the data used for the main empirical analysis. Details can be found in the data appendix C.1. The key variables of interest are measures of relative wages and relative skill supply. For relative wages I use the earnings decile ratio, i.e. the ratio of 9th-to-1st upper-earnings deciles of gross earnings for full-time dependent employees, taken from OECD [2015]. For relative skill supply, I use extrapolated educational attainment rates from Barro and Lee [2013], to calculate the ratio of the population share with at least some tertiary education over the population share with at least some primary or secondary education.

To control for the institutional environment, I use institutional data on regular and temporary employment protection, the level of wage bargaining coordination, and net replacement rates. The first two variables are index variables summarizing various items, whereby higher values indicate more restrictive institutions. Wage bargaining coordination is an index variable with higher values indicating a higher level of centralization. Net replacement rates give benefits during the first two years of unemployment as a percentage of the previous income, calculated as an average over various scenarios. Employment protection data are taken from OECD [2015]. Data on the level of wage coordination are taken from the Visser [2015]. Net replacement rates were calculated based on OECD data from ifo Institute [2013]. To

Figure 4.1: Earnings Inequality and Relative Skill Supply by Country

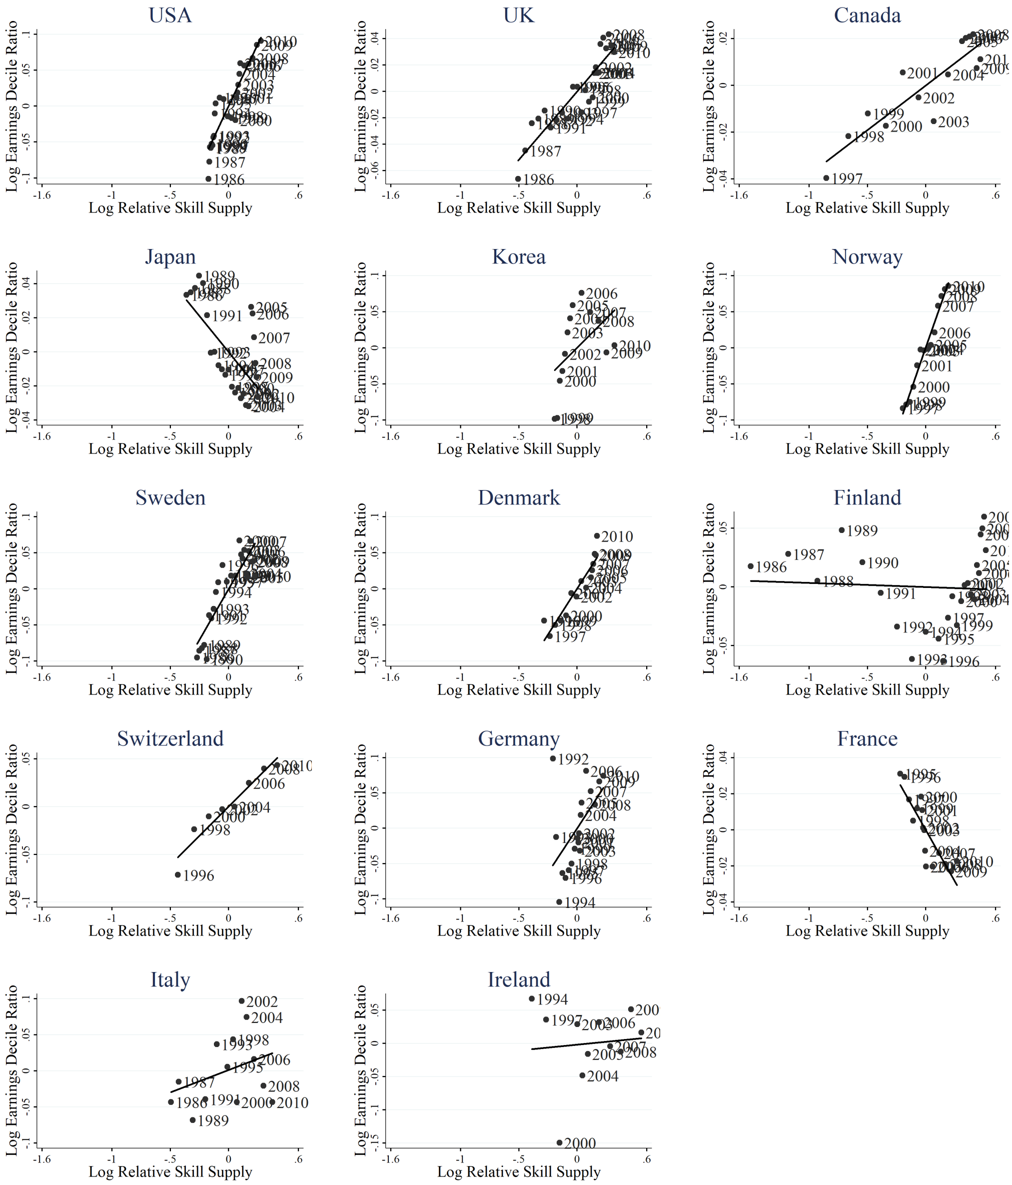

Scatterplot of log of earnings ratio plotted against the log of relative skill supply, measured as deviations from country-specific means, exemplifying the main relationship examined in this chapter. Based on annual data from 1986 to 2010. Earnings ratio is measured as the ratio of the 9th to 1st decile of gross earnings of full-time dependent employees (OECD [2015]). Relative skill supply is measured as the ratio of the percentage of the population with at least some tertiary education over the sum of the percentage of the population with at least some primary and secondary education (Barro and Lee [2013].)

reduce collinearity, I use these institutional variables to extract factor variables. Using the principal factor component analysis with oblique rotation, allowing factors to be correlated, I retrieve two factors with an Eigenvalue above one, capturing more than three fourths of the variation contained in the variables.[30] Table 4.1 reports rotated factor loadings. The first factor loads relatively heavily on regular and temporary employment protection, the second factor loads relatively heavily on wage coordination and net replacement rates and negatively on employment protection. In the following analyses I use these factor variables as measures for distinct features of the institutional environment.

Finally, as controls I use the employment to population ratio and a measure for trade openness, constructed as the sum of imports plus exports divided by gross domestic product. As robustness check, I also control for relative educational expenditures, constructed as the ratio of public and private expenditures on tertiary over primary education. All data for control variables are taken from OECD [2015], except for data on educational expenditures, which are taken from Brady et al. [2014]. Note that variables have been linearly interpolated to fill short gaps.

Overall, I have an unbalanced country-level panel dataset with 298 observations, covering the period from 1986 to 2010, and comprising 14 countries: Australia, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Korea, Norway, Sweden, Switzerland, United Kingdom, and USA.

Table 4.1: Rotated Factor Loadings for Institutional Measures

	Factor 1	Factor 2
Regular employment protection	0.8811	-0.2399
Temporary employment protection	0.7226	-0.6025
Wage coordination	0.6425	0.4474
Net replacement rate	0.5570	0.645

Rotated factor loadings for institutional variables using oblique rotation, allowing factors to be correlated. Factors with Eigenvalue larger than 1 have been extracted. Employment protection reflects index variables with higher values implying more restrictive institutions. Wage coordination reflects the level of centralization of wage negotiations. The net replacement rate gives benefits during the first two years of unemployment as a percentage of the previous income. More details for institutional variables can be found in data appendix C.1.

[30] The resulting factor variables can take on values from zero to five.

4.3.2 Extending the Katz and Murphy Approach

To relate the following analysis to the previous literature, I first extend the original approach by Katz and Murphy [1992] to a cross-country context. In particular, I regress the log of relative earnings on log relative skill supply, country-fixed effects, a common time trend and country-specific deviations from the common trend. Recall that the original approach, given by equation 4.2, suggests interpreting the estimated time trend as an estimate of average SBTC. The original approach does not control for trade openness and institutional change, so controls are omitted for now.

The hypothesis of common SBTC suggests that changes in relative earnings across countries can be explained by country-specific changes in relative skill supply and a common time trend, capturing SBTC constant across countries. Country-specific deviations from the common time trend suggest deviations from constant SBTC. I therefore run the following regression:

$$ER_{it} = \beta_0 + \beta_1 RSS_{it} + \delta_0 t + \delta_i t_{it} + \nu_i + u_{it} \tag{4.3}$$

where ER_{it} denotes the log of relative earnings for country i in period t, RSS_{it} denotes the log of relative skill supply, δ_0 is the coefficient on the common time trend and δ_i is the coefficient for country-specific deviations, ν_i denotes fixed effects, β_0 is a constant and u_{it} is the error term.[31] Note that, in all specifications to follow, I use clustered standard errors at the country level. Results are shown in table 4.2.

The first column of table 4.2 shows results for specification 4.3 when imposing $\delta_i = 0$ for all i. This is equivalent to the hypothesis of common SBTC in the approach by Katz and Murphy [1992], i.e. measuring SBTC as time trend while controlling only for relative skill supply. The trend coefficient is positive and significant, compatible with SBTC. The second column relaxes the assumption of $\delta_i = 0$, allowing for country-specific time trends. If the hypothesis of common SBTC in the Katz and Murphy framework is correct, country-specific deviations from the common trend should be insignificant. However, while the common time trend remains positive and significant, several of the country-specific deviations are significant. This analysis demonstrates that the reasoning underlying most of the literature on SBTC discussed above, i.e. assuming common SBTC while not con-

[31] Dummies for country-specific trends are constructed so that the average of time trends over all countries adds to δ_0. One can interpret δ_i as deviation from the average time trend for country i. The construction of country-specific deviations from a common time trend implies that for one country I cannot directly estimate its time trend. In particular, for Australia, the time trend is given by $\partial E_{Aus,t}/\partial t = \delta_0 - \Sigma_{i=2}^{13}\delta_i$.

Table 4.2: Common vs Country-Specific Time Trends

	(1) ER	(2) ER
RSS	-0.0612* (-2.53)	-0.0615 (-1.57)
Trend	0.00596** (3.93)	0.00827*** (18.63)
Trend Can		0.00178 (0.50)
Trend Den		0.00289** (3.28)
Trend Fin		-0.00259 (-1.14)
Trend Fra		-0.0101*** (-15.37)
Trend Ger		-0.000967* (-2.50)
Trend Ire		-0.00537** (-3.05)
Trend Ita		-0.00426*** (-5.36)
Trend Jap		-0.00869*** (-14.97)
Trend Kor		0.00443** (3.75)
Trend Nrw		0.00747*** (10.19)
Trend Swe		-0.000787 (-1.82)
Trend Swz		0.00255 (1.35)
Trend UK		-0.00272** (-3.26)
Trend USA		-0.000206 (-0.71)
Fixed Effects	Yes	Yes
Period	1986-2010	1986-2010
N	298	298

Note: t statistics in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Results for regression specification 4.3. Dependent variable 'ER' is log of 9th to 1st earnings decile ratio. 'RSS' is log of relative skill supply. 'Trend' denotes common trend. Deviations from common trend are given by country-suffix trend. Using clustered standard errors at country-level. Regression includes constant. Trend deviation for Australia not shown. As argued for in the text, country-specific deviations from the common trend suggest variation in skill-bias.

trolling for institutional change, is rejected in this simple analysis. Results instead suggest substantial variation across countries. Of course, country-specific deviations might reflect changes in country-specific factors, e.g. trade openness or institutional change. I address this concern in the following subsection.

4.3.3 Country-Specific Skill Bias

Next, I examine SBTC variation in terms of the slope between measures for relative wages and relative skill supply, adding additional control variables. To do this, I estimate the following estimation specification:

$$ER_{it} = \beta_0 + \beta_1 RSS_{it} + \gamma_i RSS_{it} + \boldsymbol{\theta' X_{it}} + \nu_i + u_{it} \tag{4.4}$$

where terms are as before and $\boldsymbol{X_{it}}$ is a vector of control variables, including a measure of trade openness, the employment to population ratio, and institutional measures. Country-specific interaction dummies γ_i are constructed such that β_1 measures the common slope and γ_i indicates country-specific deviations.[32] For specification 4.4, the slope identifying SBTC is given by the marginal effect of relative skill supply

$$\frac{\partial ER_{it}}{\partial RSS_{it}} = \beta_1 + \gamma_i$$

This specification allows decomposing the slope into a common and a country-specific component. The hypothesis of common SBTC can be tested as follows: if countries face common SBTC, the same positive relationship applies to all countries, so $\beta_1 > 0$ and $\gamma_i = 0$ for all i. Conversely, if countries exhibit SBTC with varying magnitude, we have $\beta_1 > 0$ and $\gamma_i \neq 0$ for some i.

Results are reported in table 4.3. The first column omits all control variables except for country-fixed effects. This specification amounts to estimating whether differences shown graphically in figure 4.1 are statistically significant. Table 4.3 confirms that it is: the common slope is positive, suggesting SBTC, and all coefficients on country-specific deviations are significant. However, as argued above, finding country-specific slope effects may simply reflect country-specific variation in confounding variables. Columns two to five successively include further control variables.

First, I add trade openness and the employment to population ratio as basic controls. Trade openness has been considered as a competing explanation in the literature. Also, in a cross-sectional context, controlling for relative skill supply by educational attainment suggest keeping the employment to population ratio constant to account for differences in unemployment and labour force participation.[33] Results are reported in column two of table 4.3. The inclusion of controls does not affect the average effect very much but reduces the significance of country-specific deviations. However, deviations generally remain significant. Deviations only for Germany and Korea turn insignificant.

Second, I additionally control for institutional variables, which have been mentioned in the literature as competing explanation for changes in relative wages.

[32] In particular, γ_i is constructed as $\gamma_i = \nu_i - \nu_1$ for $i = 2, \ldots, n$ so we have $\gamma_j = 1$ for observations on j and $\gamma_j = -1$ for observations on $i = 1$. This assures that one can interpret β_1 as the average slope coefficient, and γ_i as deviations thereof. For $i = 1$, which refers to Australia, the slope coefficients is given by $\partial ER_{1t}/\partial RSS_{1t} = \beta_1 - \Sigma_{i=1}^{n} \gamma_i$.

[33] Countries with similar relative skill supply may differ in their effective supply if some skilled or unskilled workers are unemployed or do not participate in the labour market. Ideally, one would want to control for participation rates by educational endowment. The overall participation rate serves as proxy.

Specifications in columns three to six of table 4.3 successively add factor variables and their interaction. Recall that I use factor variables to proxy for the institutional environment. I use factor variables, rather than individual institutional measures, to account for institutional complementarities and reduce collinearity.[34] Adding institutional measures, by themselves and including their interaction, acts to reduce the common slope, and renders some country-specific deviations from the common slope insignificant. In particular, this is the case for Sweden, Switzerland and the UK. However, institutional measures are not themselves significant. Country-specific deviations from common SBTC remain significant even after controlling for confounding variables. These results continue to reject the hypothesis that SBTC is constant across countries.

Given that results suggest SBTC differs across countries, one might ask which country-level factors can explain this variation. I address this question next. In particular, I ask whether cross-country SBTC differences are systematically related to institutional variation. Recall that the above literature review suggests some arguments that institutions affect technology, or SBTC in particular. Also, recall that some authors covered in the review suggest that more rigid institutions, introducing a binding lower bound for unskilled workers, exhibit lower SBTC.

[34] Note that more institutions than those included for the extraction of factor variables are mentioned in the literature, e.g. minimum wages and union coverage, but in principle also institutional measures for financial and legal systems. I use a restricted set of institutions because of limited data coverage. The fact that I do not cover all relevant institutional variables implies that results may still suffer from omitted variable bias.

Table 4.3: Country-Specific Effects

	(1) ER	(2) ER	(3) ER	(4) ER	(5) ER
RSS	0.146*** (7.12e+11)	0.148*** (4.48)	0.133** (3.99)	0.115** (3.54)	0.115** (3.51)
RSS Can	-0.108*** (-4.61e+11)	-0.123** (-3.03)	-0.113** (-2.99)	-0.103* (-2.95)	-0.105** (-3.18)
RSS Den	0.106*** (3.93e+11)	0.108*** (15.61)	0.131*** (6.64)	0.154*** (7.78)	0.154*** (7.08)
RSS Fin	-0.150*** (-7.22e+11)	-0.140*** (-5.16)	-0.126*** (-4.40)	-0.112** (-3.86)	-0.112*** (-4.20)
RSS Fra	-0.260*** (-1.83e+11)	-0.279*** (-11.12)	-0.265*** (-10.20)	-0.271*** (-10.05)	-0.267*** (-8.60)
RSS Ger	0.107*** (5.49e+10)	0.124 (0.95)	0.0807 (0.64)	0.0140 (0.11)	0.0177 (0.15)
RSS Ire	-0.135*** (-4.12e+11)	-0.184* (-2.39)	-0.193* (-2.73)	-0.202** (-3.15)	-0.206** (-3.09)
RSS Ita	-0.0767*** (-3.73e+11)	-0.0850** (-3.42)	-0.124** (-3.16)	-0.164** (-4.10)	-0.144 (-2.06)
RSS Jap	-0.231*** (-1.12e+12)	-0.229*** (-38.12)	-0.255*** (-12.16)	-0.259*** (-13.30)	-0.257*** (-17.26)
RSS Kor	0.0160*** (2.14e+10)	0.0321 (0.55)	0.0613 (0.97)	0.0840 (1.41)	0.0837 (1.41)
RSS Nrw	0.316*** (2.68e+11)	0.322*** (6.57)	0.360*** (6.52)	0.451*** (6.42)	0.446*** (7.56)
RSS Swe	0.135*** (3.16e+11)	0.183** (3.35)	0.128 (1.95)	0.0740 (1.17)	0.0754 (1.22)
RSS Swz	-0.0261*** (-1.27e+11)	-0.0213** (-3.86)	0.000142 (0.01)	0.0267 (1.32)	0.0279 (1.20)
RSS UK	-0.0427*** (-2.08e+11)	-0.0526* (-2.64)	-0.0339 (-1.54)	-0.0286 (-1.43)	-0.0336 (-1.66)
RSS USA	0.197*** (4.74e+11)	0.211*** (7.04)	0.226*** (7.14)	0.204*** (5.90)	0.187* (2.72)
Inst 1			-0.0313 (-1.23)	-0.0360 (-1.53)	-0.0260 (-0.72)
Inst 2				0.0180 (1.68)	0.0291 (0.55)
Inst 1*2					-0.00377 (-0.22)
Emp/Pop		0.00180 (0.48)	0.00213 (0.57)	0.00178 (0.48)	0.00168 (0.47)
Trade		-0.0508 (-0.38)	-0.0786 (-0.60)	-0.0824 (-0.68)	-0.0804 (-0.66)
Fixed Effects	Yes	Yes	Yes	Yes	Yes
Period	1986-2010	1986-2010	1986-2010	1986-2010	1986-2010
N	298	298	298	298	298

Note: t statistics in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Results for regression specification 4.4. The dependent variable 'ER' is the log of the 9th to 1st earnings decile. 'RSS' is the log of relative skill supply and measures the common slope. Deviations from common trend are given by country-suffix RSS. Using clustered standard errors at country-level. Regression includes constant. Slope deviation for Australia not shown. As argued for in the text, SBTC is given by the marginal product of RSS. Country-specific deviations from the common slope indicate SBTC differences.

4.3.4 Institution-Specific Skill Bias

Having found significant SBTC differences across countries, I ask in this subsection whether differences are related to institutions. Recall that I identify SBTC in terms of the slope of an increase in log relative skill supply on log relative earnings. To test whether slopes vary systematically with institutional measures, I replace country-specific interactions with relative skill supply by interactions between institutional measures and relative skill supply. In particular, I estimate the following specification:

$$ER_{it} = \beta_0 + \beta_1 RSS_{it} + \boldsymbol{\beta_2 Inst_{it}} + \boldsymbol{\beta_3 Inst_{it}} * RSS_{it} + \boldsymbol{\theta' X_{it}} + \nu_i + u_{it} \quad (4.5)$$

Terms are as before. I use three sets of institutional measures.[35] This specification decomposes the slope into a common and an institution-specfic component:

$$\frac{\partial ER_{it}}{\partial RSS_{it}} = \beta_1 + \boldsymbol{\beta_3 Inst_{it}}$$

Because the slope indicates SBTC, this specification effectively decomposes SBTC into a common and an institution-specific component. The previous analysis established that SBTC varies across countries. If SBTC differences are related to institutions, we would expect $\boldsymbol{\beta_3} \neq 0$, whereas we would expect $\boldsymbol{\beta_3} = 0$ if SBTC differences are unrelated to institutions. If, as is suggested by Acemoglu [2003], Pischke [2005], and Fadinger and Mayr [2015], more rigid institutions lead to smaller SBTC, we would expect $\boldsymbol{\beta_3} < 0$. That is, more rigid institutions are associated with a shallower slope, or a negative one for sufficiently rigid institutions. Note that, because I control for country-fixed effects, institutions are effectively measured as deviations from country-specific means. The effect of institutional variation is effectively identified by countries with decreasing or increasing institutional rigidity exhibiting smaller or larger SBTC.

Results are shown table 4.4. Columns one to three successively add institutional variables. There is only mixed evidence for slope differences being systematically related to country-specific changes in institutions. Including factor variables by themselves does not yield significant effects, although coefficient signs on institutional measures suggest that SBTC is shallower for more rigid institutions. When including the full set of institutional measures, the interaction term is significant. The SBTC measure, the marginal effect of an increase in log relative skill supply,

[35] I include only the first factor variable, both factor variables, and both factor variables and their interaction.

depends on the level of institutional measures. To see how SBTC varies with institutions when interactions are included, figure 4.2 depicts the marginal effect of an increase in log relative skill supply for the specification in column three, and the corresponding confidence intervals at the ten percent significance level, for different institutional values. To help relating this graph to particular countries covered in the dataset, figure C.1 plots the country-specific mean for the factor variables, and depicts institutional variation as circles around the respective mean, with circle size proportional to the sum of variance for both factors.

Starting from the left, subsequent boxes of figure 4.2 show the marginal effect of an increase in relative skill supply at different values of factor 1 and factor 2. The marginal effect indicates the partial correlation, or the slope, interpreted as SBTC. In this light, the marginal effect is insignificant for low values of factors 1 and 2. For low values of factor 1, SBTC becomes stronger, and ultimately significant at the ten percent level, for higher values of factor 2. For given values of factor 2, SBTC becomes weaker for higher values of factor 1, and significantly so for sufficiently high values of factor 1. This provides some evidence for a systematic relationship between measures for distinct features of the institutional environment and SBTC. Recall that factor 1 loads relatively heavily on employment protection, and factor 2 loads relatively heavily on wage bargaining coordination and net replacement rates. Interpreting institutional measures in more substantive terms, this suggests that countries with low levels of employment protection have higher SBTC the higher the level of wage bargaining and net replacement rates. On the other hand, for a given level of wage bargaining and net replacement rates, higher employment protection implies lower SBTC. For example, figure C.1 shows that Australia and Sweden exhibit similar levels of factor 2, measuring intermediate levels of wage bargaining coordination and net replacement rates, but Australia has low and Sweden has high employment protection. Results would suggest that SBTC is higher in Australia than in Sweden. Somewhat counterintuitively, however, results also suggest that countries with both low employment protection and wage bargaining levels and net replacement rates, such as the UK and US, exhibit no SBTC. These latter results may be explained by the limited institutional variation exhibited by countries with low factor values, as shown in figure C.1, making it difficult to identify any effects. Overall, these results provide some tentative evidence that institutional measures may be systematically related to SBTC differences, and that *some* institutional measures are related to lower SBTC. If so, one may ask next whether SBTC differences found before relate to country-specific factors other than the institutional measures included, or whether they can be fully explained by institutional measures. I address

Figure 4.2: Institution-Specific Marginal Effects with Interaction

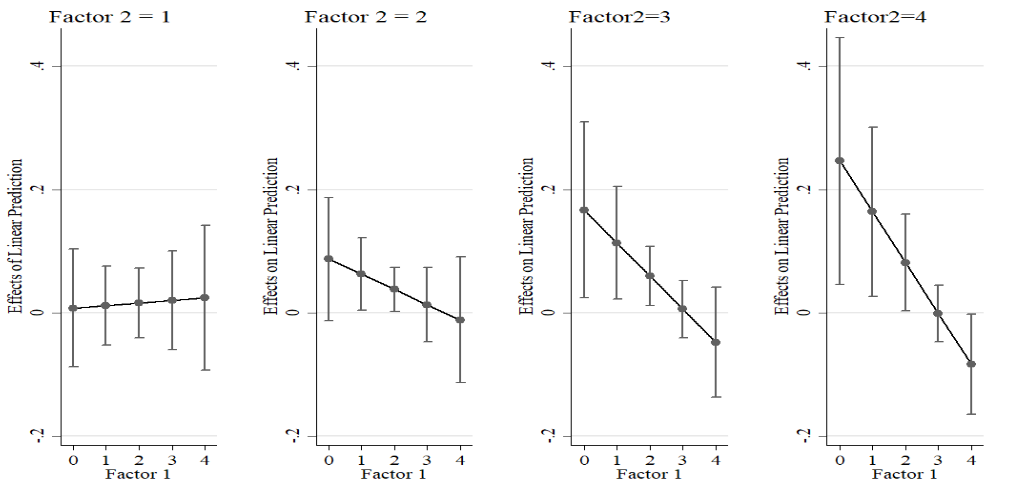

Marginal effects at the ten percent significance level for specification 4.5, for results shown in column 3 of table 4.4. Marginal effects are shown for selected values of factor variables 1 and 2. Factor variables range from 0 to 5. Factor 1 loads heavily on employment protection, factor 2 loads heavily on wage bargaining coordination and net replacement rates. Higher values indicate more restrictive institutions. Details for factor variables are given in 4.3.1. As argued for in the text, marginal effects indicate SBTC. The figure shows how SBTC varies across countries with institutional measures.

this question in the next analysis.

4.3.5 Country- Versus Institution-Specific Skill Bias

To test whether country-specific SBTC variation is explained by institutional measures or other country-specific factors, I augment the previous specification by adding country-specific effects. I estmate the specification:

$$ER_{it} = \beta_0+\beta_1 RSS_{it}+\gamma_i RSS_{it}+\boldsymbol{\beta_2' Inst_{it}}+\boldsymbol{\beta_3' Inst_{it}}*RSS_{it}+\boldsymbol{\theta' X_{it}}+\nu_i+u_{it} \quad (4.6)$$

This model nests the hypothesis for institution-specific differences only for $\gamma_i = 0$ for all i, and for country-specific differences only for $\boldsymbol{\beta_3} = \mathbf{0}$. The slope for the relationship between relative wages and relative skill supply is given by

$$\frac{\partial ER_{it}}{\partial RSS_{it}} = \beta_1 + \gamma_i + \boldsymbol{\beta_3' Inst_{it}}$$

Terms are as before. This analysis decomposes the slope into a common effect, a country-specific effect, and an institution-specific effect. Note that the country-specific effect comprises all time-invariant variables, including time-invariant institutional variables, whereas the institution-specific effect measures the effect for time-variant institutions covered by the sample. If country-specific SBTC differences are explained by institution-specific effects only, we would expect $\gamma_i = 0$ for all i. If SBTC differences additionally depend on other country-specific time-invariant factors, we would expect $\beta_3 \neq 0$ and $\gamma_i \neq 0$ for some i.

Results are shown in table 4.5. Columns one and two repeat the basic model, excluding basic and/or institutional controls. Columns three to five successively add institution-specific interactions. Overall, country-specific effects remain significant, although the significance is generally reduced. There is some evidence that time-variant institutions are explaining SBTC variation in addition to country-specific effects: coefficients for the institutional measure factor 1 are negative and significant if included as single variables. To examine the slope for the specification including an interaction between both institutional measures, figure 4.3 gives the marginal effect, at the ten percent significance level, of an increase in log relative skill supply for selected values of both factor variables.[36] The same pattern emerges as before: for given values of factor 2, measuring mostly wage bargaining coordination and net replacement rates, a higher value of factor 1, measuring mostly employment protection, is associated with smaller SBTC, or no SBTC at all. For low values of factor 1, higher values of factor 2 imply larger SBTC. Again, however, countries with generally low levels of institutional rigidities are implied to suggest no SBTC. Results are generally more significant compared to specification 4.5, which does not include country-specific effects. Overall, these results suggest that time-variant institutional measures can explain some SBTC differences, but other country-specific factors remain important.

[36]Marginal effects are shown for average country-specific effects.

Table 4.4: Institution-Specific Effects

	(1) ER	(2) ER	(3) ER
RSS	0.116*	0.121*	-0.0719
	(2.21)	(2.18)	(-0.95)
Inst 1	-0.0823	-0.0745	0.0710
	(-1.40)	(-1.36)	(0.74)
Inst 2		0.00765	0.174**
		(0.32)	(3.29)
Inst 1*2			-0.0626**
			(-3.01)
RSS Inst 1	-0.0315	-0.0319	0.0331
	(-1.38)	(-1.51)	(1.17)
RSS Inst 2		-0.00272	0.0795
		(-0.25)	(2.03)
RSS Inst 1*2			-0.0289*
			(-2.97)
Emp/Pop	0.000893	0.000364	0.000126
	(0.26)	(0.10)	(0.04)
Trade	0.0800	0.0811	0.131
	(0.93)	(0.98)	(1.68)
Fixed Effects	Yes	Yes	Yes
Period	1986-2010	1986-2010	1986-2010
N	298	298	298

Note: t statistics in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Results for regression specification 4.5. The dependent variable 'ER' is the log of the 9th to 1st earnings decile. 'RSS' is the log of relative skill supply and measures the common slope. 'Inst 1' and 'Inst 2' refer to factor variables 1 and 2, extracted from institutional variables. Details for factor variables given in 4.3.1. Using clustered standard errors at country-level. Regression includes constant. Slope deviation for Australia not shown. As argued for in the text, marginal effects indicate SBTC. Regression results show how SBTC varies across countries with institutional measures.

Table 4.5: Country- & Institution-Specific Effects

	(1) ER	(2) ER	(3) ER	(4) ER	(5) ER
RSS	0.146*** (7.12e+11)	0.148*** (4.48)	0.318** (3.64)	0.331*** (5.07)	-0.127 (-0.51)
RSS Can	-0.108*** (-4.61e+11)	-0.123** (-3.03)	-0.218** (-3.41)	-0.214** (-3.96)	-0.0962 (-1.29)
RSS Den	0.106*** (3.93e+11)	0.108*** (15.61)	0.159*** (9.14)	0.218*** (5.62)	0.171* (2.95)
RSS Fin	-0.150*** (-7.22e+11)	-0.140*** (-5.16)	-0.0710 (-1.96)	-0.0416 (-1.54)	-0.0237 (-0.92)
RSS Fra	-0.260*** (-1.83e+11)	-0.279*** (-11.12)	-0.196*** (-5.67)	-0.210*** (-5.51)	-0.170*** (-4.32)
RSS Ger	0.107*** (5.49e+10)	0.124 (0.95)	0.115 (1.03)	0.0426 (0.35)	0.107 (0.88)
RSS Ire	-0.135*** (-4.12e+11)	-0.184* (-2.39)	-0.273** (-3.08)	-0.248* (-2.48)	-0.340* (-2.79)
RSS Ita	-0.0767*** (-3.73e+11)	-0.0850** (-3.42)	0.0754 (0.92)	-0.0882 (-0.45)	-0.263* (-2.82)
RSS Jap	-0.231*** (-1.12e+12)	-0.229*** (-38.12)	-0.312*** (-10.75)	-0.316*** (-15.93)	-0.283*** (-11.97)
RSS Kor	0.0160*** (2.14e+10)	0.0321 (0.55)	0.127 (1.71)	0.135 (1.79)	0.158 (1.93)
RSS Nrw	0.316*** (2.68e+11)	0.322*** (6.57)	0.473*** (7.07)	0.561*** (7.98)	0.538*** (10.21)
RSS Swe	0.135*** (3.16e+11)	0.183** (3.35)	0.135* (2.38)	0.0930 (1.60)	0.151* (2.71)
RSS Swz	-0.0261*** (-1.27e+11)	-0.0213** (-3.86)	-0.0104 (-1.83)	0.0250 (1.18)	0.0233 (1.15)
RSS UK	-0.0427*** (-2.08e+11)	-0.0526* (-2.64)	-0.120* (-2.76)	-0.123** (-3.77)	0.0101 (0.14)
RSS USA	0.197*** (4.74e+11)	0.211*** (7.04)	0.0852 (1.41)	0.119 (1.19)	-0.0265 (-0.20)
Inst 1			-0.169* (-2.53)	-0.165* (-2.97)	0.00967 (0.10)
Inst 2				-0.0134 (-0.35)	0.219 (1.54)
Inst 1*2					-0.0728 (-1.94)
RSS Inst 1			-0.0765* (-2.26)	-0.0680 (-2.09)	0.0577 (1.00)
RSS Inst 2				-0.0212 (-0.85)	0.155 (1.49)
RSS Inst 1*2					-0.0516 (-2.05)
Emp/Pop		0.00180 (0.48)	0.00293 (0.85)	0.00293 (0.87)	0.00384 (1.25)
Trade		-0.0508 (-0.38)	-0.109 (-0.82)	-0.121 (-0.98)	-0.108 (-0.86)
Fixed Effects	Yes	Yes	Yes	Yes	Yes
Period	1986-2010	1986-2010	1986-2010	1986-2010	1986-2010
N	298	298	298	298	298

Note: t statistics in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Results for regression specification 4.6. The dependent variable 'ER' is the log of the 9th to 1st earnings decile. 'RSS' is the log of relative skill supply and measures the common slope. Deviations from common trend are given by country-suffix RSS. Inst 1 and 2 refer to factor variables 1 and 2, extracted from institutional variables. Details for factor variables given in 4.3.1. Using clustered standard errors at country-level. Regression includes constant. Slope deviation for Australia not shown. As argued for in the text, marginal effects indicate SBTC. Regression results show how SBTC varies with country-specific time-invariant factors and time-variant institutional measures.

4.3.6 Institutional Change Versus Institutional Differences

The above results suggest that SBTC differs among the OECD countries included in the sample, and that some of these differences are related to institutional measures. As mentioned above, the effect of institutional measures is identified in terms of country-specific institutional variation: I identify the level effect of institutional measures by estimating whether relative earnings and relative skill supply are more (less) positively (negatively) correlated if institutions are comparatively rigid relative to their country-specific mean. One may wonder whether effects for institutional measures identified in terms of institutional variation within countries are similar to effects identified in terms of institutional differences across countries. Identifying institutional differences in the cross-section implies that one cannot control for country-fixed effects. As an intermediate solution, I use the following method, also known as fixed effects vector decomposition, focusing on institution-specific effects only:[37] in the first stage, I retrieve estimates for fixed effects and, in the second stage, orthogonalize these with respect to institutional variables. Then, I estimate specification 4.5 while including orthogonalized fixed effects in the third stage. Thus, I effectively use the within-estimator for all non-orthogonalized variables, i.e. the basic controls and relative skill supply, while using the random effects estimator for institutional variables. The aim is to estimate coefficients more efficiently while exploiting between-country variation in levels. Coefficients for institutional variables are a combination of the within- and between-estimator, and thus reflect the variation in institutional variables across countries in levels. In contrast to previous specifications, one limitation of this approach is that coefficient estimates for institutional measures may suffer from omitted variable bias for any time-invariant factors. With this in mind, results for the within- and between-estimators as well as the second and third stage are reported in table 4.6. Overall, results suggests that differences in SBTC are related to institutional measures identified between countries. Figure 4.4 depicts marginal effects, as before at the ten percent significance level, for selected values of each factor variable. Marginal effects confirm the pattern observed before, but results are much more significant: for low values of factor 1, higher values of factor 2 imply larger SBTC. For given values of factor 2, SBTC is higher the lower is the value of factor 1. Countries with low values for both factors exhibit no significant SBTC, however. Figure C.1 implies that, based on their institutional environment only, countries such as the US or UK exhibit no SBTC. As institution-specific effects are identified in the cross-section, this cannot

[37] See Pluemper and Troeger [2007].

Figure 4.3: Country & Institution-Specific Marginal Effects with Interaction

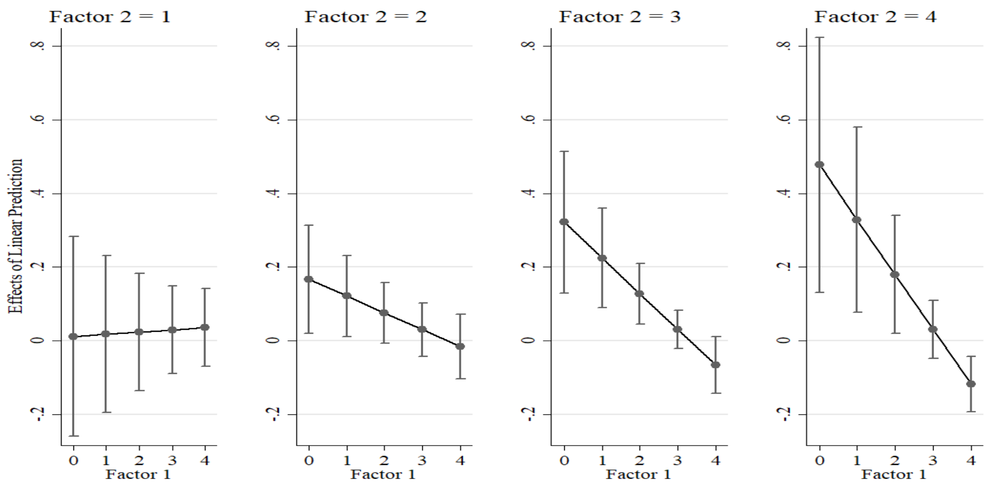

Marginal effects at the ten percent significance level for specification 4.6. Results shown in column 6 of table 4.5. Marginal effects shown for selected values of factor variables 1 and 2 at average country-specific effects. Factor variables range from 0 to 5. Higher values indicate more restrictive institutions. Details for factor variables are given in 4.3.1. As argued for in the text, marginal effects indicate SBTC. The figure shows how SBTC varies across countries with institutional measures for average country-specific effects.

be explained by low institutional variation within these countries.

4.3.7 Robustness Checks

In this section, I discuss robustness checks. In particular, I first check whether results are driven by any particular country by dropping one country at a time and repeating the analysis. Second, I check whether results are sensitive to the measure for wage inequality, and, third, I control for educational quality.

First, I discuss casewise dropping of countries. To see whether results are driven by any single country, I reestimate the main estimation specifications while dropping one country at a time. In particular, I estimate the full specifications 4.4, 4.5, and 4.6, i.e. including all institutional measures and the interaction. Results are not reported here. However, omitting any single country does not change qualitative conclusions for the main analysis. Results are not driven by any single country.

Second, I turn to alternative wage inequality measures. To best fit the difference between wages for skilled and unskilled workers, I chose as a measure for wage inequality the earnings decile ratio for the 9th to 1st decile. However, it has been argued that the general trend for increasing wage inequality, as captured by this

Table 4.6: Identifying Institutional Effects Within and Between Countries

	Within Est.	Between Est.	Random Effects	2nd Stage	3rd Stage
RSS	-0.108	2.072	0.0919		-0.108*
	(-1.16)	(1.14)	(0.66)		(-2.44)
Inst 1	0.00321	-0.911	-0.247**	-0.179***	-0.176***
	(0.03)	(-1.25)	(-4.00)	(-3.76)	(-14.26)
Inst 2	0.192**	-2.109	-0.0879	-0.132**	0.0598
	(3.22)	(-1.20)	(-0.68)	(-2.81)	(1.72)
Inst 1*2	-0.0550*	0.754	0.0368	0.0281	-0.0269**
	(-2.68)	(1.05)	(0.88)	(1.29)	(-2.81)
RSS Inst 1	0.0175	-0.582	-0.0142	0.0123	0.0298***
	(0.62)	(-0.89)	(-0.43)	(0.85)	(4.02)
RSS Inst 2	0.118*	-1.589	0.0521	0.0994***	0.218***
	(2.23)	(-1.10)	(0.51)	(4.88)	(6.51)
RSS Inst 1*2	-0.0349*	0.574	-0.0141	-0.0299***	-0.0648***
	(-2.98)	(0.96)	(-0.48)	(-3.52)	(-7.53)
Emp/Pop	0.00272	-0.0395	-0.00647		0.00272
	(0.86)	(-1.68)	(-1.18)		(1.87)
Trade	0.184*	0.567	-0.200		0.184***
	(2.56)	(0.88)	(-1.70)		(12.85)
FE	Yes	No	No	No	Yes/No
N	254	12	254	254	254

Note: t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Results for regression specification 4.5 using fixed effects vector decomposition. Third stage contains fixed effects orthogonalized for factor variables. Dependent variable 'ER' is the log of the 9th to 1st earnings decile. 'RSS' is the log of relative skill supply and measures the common slope. Inst 1 and 2 refer to factor variables 1 and 2, extracted from institutional variables. Details for factor variables given in 4.3.1. Using clustered standard errors at country-level. Regression includes constant. Slope deviation for Australia not shown. As argued for in the text, marginal effects indicate SBTC. Regression results shows how SBTC varies across countries with institutional measures when identified in levels.

ratio, would mask important changes in wage structure dynamics: while in earlier decades increasing wage inequality would have affected the entire wage distribution, it would now be concentrated in the lower or upper half.[38] This suggests that results may affect different parts of the earnings distribution differently. It is therefore instructive to repeat the analysis using wage inequality measures for different parts of the wage distribution. I thus reestimate the main estimation specifications 4.4, 4.5, and 4.6, while replacing the 9th to 1st earnings decile ratio with a) the 9th to 5th earnings decile ratio and b) the 5th to 1st earnings decile ratio. Results are reported in table C.1. Overall, results across all three specifications suggest that the conclusions in the main analysis are driven by changes in the upper half rather than the lower half of the wage distribution, although SBTC differences aross countries are robustly observed for each dependent variable. Interestingly, coefficients on in-

[38] See Lemieux [2007].

Figure 4.4: Institution-Specific Marginal Effects Identified in Levels

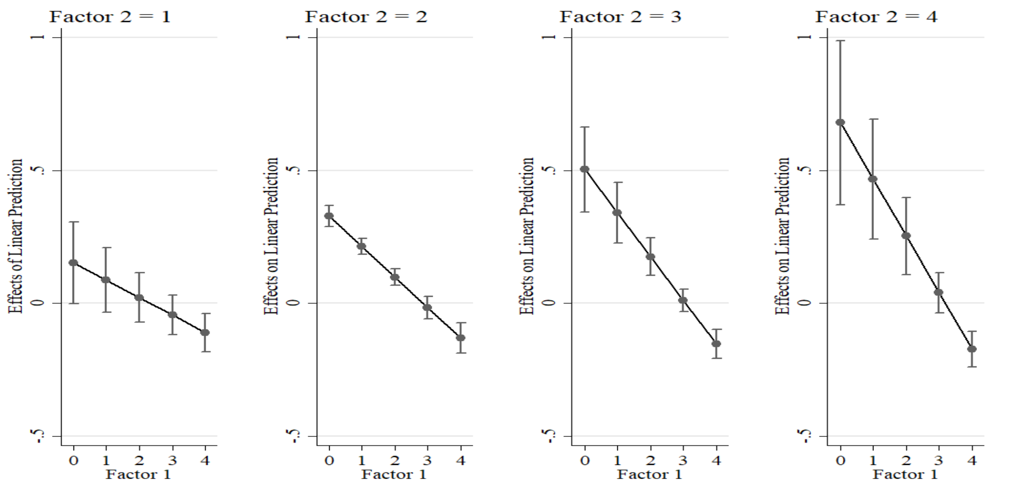

Marginal effects at the ten percent significance level for specification 4.5 for the third stage of the fixed effects vector decomposition. Results shown in column 6 of table 4.6. Marginal effects shown for selected values of factor variables 1 and 2. Factor variables range from 0 to 5. Higher values indicate more restrictive institutions. Details for factor variables are given in 4.3.1. As argued for in the text, marginal effects indicate SBTC. The figure shows how SBTC varies across countries with institutional measures when identified in levels.

stitutional measures become more significant when using the earnings ratio for the upper half rater than the lower half. Based on the review of theoretical arguments on SBTC variation, one may have expected that institutions affect SBTC to the extent they are binding for unskilled workers. Instead, it appears that, if institutions are important in accounting for SBTC variation, this may be due to institutions affecting the upper or the entire wage distribution.

Third, I additionally control for differences in educational quality. One may argue that slope variation is driven by changes in educational quality across countries: consider two countries with a similar increase in relative skill supply. Ceteris paribus, if one country improves its educational quality for tertiary education relative to the other country over this period, its measured increase in relative skill supply would be understated, and different changes in relative earnings may be misconceived as differences in SBTC. To control for such effects, I additionally include a measure for relative educational quality, constructed as the ratio of public and private expenditure on tertiary over primary education. Note that these data do not cover Canada, Korea, Japan, Sweden, Norway and Switzerland. I thus compare the baseline model with the extended model including the additional controls for only a subset of the sample, covering the same observations. Results for the main estimation specifications with either set of control variables are reported in table C.2. Educational quality appears to be significantly related to relative wages: surprisingly, higher relative educational quality implies lower relative wages. The conclusions in the main analysis are not affected: SBTC exhibits significant differences across countries. Interestingly, coefficients on institutional measures become more significant in the baseline model and the model using additional controls. This may reflect the omission of countries with little institutional variation. Overall, this suggests that the findings in the main analysis are not due to omitting educational quality controls.

4.3.8 Summary

The empirical analysis showed that there is no evidence in favour of common SBTC among OECD countries. In contrast, there is strong evidence for SBTC varation for the countries included in the sample, which cover a wide range of diverse OECD countries. There is some evidence for this variation to be related to time-variant institutions. However, it is notoriously difficult to identify the effect of relative time-invariant institutions when controlling for fixed effects. There is some evidence for institutions in levels being related to SBTC when omitting fixed effects, but these may reflect omitted variables. Results seem to be driven by changes in the upper half rather than the lower half of the wage distribution. These conclusions are not driven by single countries, and are robust to controling for educational quality.

4.4 Discussion

As set out in the introduction, the widespread assumption of common SBTC has important implications about the scope of policy makers to affect the impact of technological change on the labour market, but, despite this importance, rests on ad hoc reasoning rather than empirical evidence. At the same time, however, assessing SBTC across countries faces many limiations: because technology is not observed, SBTC is generally identified in residual terms for changes in relative labour market outcomes, many confounding factors need to be controlled for, but data coverage is limited, and measurement errors are severe. These limitations apply to identifying SBTC within countries, but may be even more serious when comparing SBTC across countries. I therefore adopted a rethorical approach to test whether SBTC varies across OECD countries, and whether such variation is related to labour market institutions. I extended the approach by Katz and Murphy [1992] to a cross-country context, additionally controlling for relevant confounding variables not systematically included in the literature. The reasoning is that, if one accepts findings from the seminal literature establishing the presence of SBTC, one ought to accept conclusions from this approach about SBTC differences.

Using this approach, I tested whether SBTC varies across countries. Findings suggest differences in SBTC across countries which are robust to including competing explanations, trade openness and institutions, in a unified framework. Data on institutions is scarce, however, so I used factor variables derived from several labour market institutions to proxy for the institutional environment. This leaves open the possibility that differences across countries reflect changes in institutions not captured by these proxies. Assuming the validity of the approach by Katz and Murphy [1992], results suggests that countries differ in SBTC, and that country-specific factors, presumably under the control of policy makers, affect the magnitude of the impact of technological changes on the labour market. This provides evidence that the assumption of common SBTC is unwarranted. There are reasons to believe that SBTC is not exogenous to the economy.

Next, I tested whether these country-level factors reflect institutional variation across countries. The review of theoretical arguments on country-level determinants of technology in fact suggested that institutions do affect technology in general, and SBTC in particular may be dampenend by more rigid institutions. Findings provide some evidence that institutional measures vary systematically with SBTC, although other country-specific factors appear to be important as well. There is only mixed evidence, however, that more rigid institutions imply lower SBTC. Results suggest this is the case for some institutions but not for others. Interpreting institutional measures in more substantive terms, this is arguably the case for employment protection but not wage coordination or net replacement rates. Also, the estimated relationship between SBTC and institutional measures provides some counterintuitive results, which are not in line with arguments that more rigid institutions, introducing wage floors for unskilled workers, imply lower SBTC: the most liberal countries, e.g. the US and UK, would exhibit no SBTC on this account. While this may reflect difficulties in identifying institutional effects when controlling for country-fixed effects, these conclusions are not altered much when identifying institutional effects using between-country variation.

Overall, findings in this chapter provide evidence against the pessimistic view that the scope of policy makers in controlling the impact of technological change on the labour market is limited to reacting to exogenously given technological shocks. SBTC appears to vary with country-specific factors, which arguably are under some control of policy makers. It appears intuitive that institutions comprise part of these country-specific factors. Given the difficulty to identify institutional effects in this context, however, findings are merely suggestive. Further research is required to examine the link between institutions, or other country-specific factors, and technological change.

Appendix A

(For Chapter 2)

A.1 Data

A.1.1 General Remarks

ASHE exhibits discontinuities in 2004 and 2006. In 2004 supplementary information was added, and in 2006 automatic occupational coding was introduced. For each year, two datasets are available, of which I use the one incorporating the most recent methodology. Note that the discontinuity in 2004 does not seem to affect variables used for this analysis.

Both NESPD and ASHE feature a small number of duplicate observations, i.e. several observations for the same worker in the same year. To obtain a panel dataset at the worker-level, I only retain observation on each worker's main job. In most cases duplicates reflect multiple jobs, in which case I retain the observation associated with highest annual gross pay (agp) in AHSE and annual gross pay (or in case of missing values for the former gross weekly pay (gpay), or gross weekly pay excluding overtime (gpox)) in NESPD.

A small number of observations exhibit inconsistencies or missing values in occupation, age, or gender variables. I deal with these inconsistencies as follows. I replace missing occupational codes (less than 0.01 percent of all observations) with the worker's last non-missing value. In a very small number of cases I recode the occupational variable due to an apparent typo in the original dataset (SOC 2010 code 232, which does not exist, to 223, which is the workers previous and past SOC code). Also, in a very small number of cases there are missing values for the variable age, in which case I fill in values consistent with the previously or subsequently reported age. Only in one case no age can be imputed, so I delete this observation. About 3 percent of observations feature inconsistent information on age or gender, i.e. the age reported from one year to the next differs by more than 2 years or gender changes from one year to the other. In the first case I extrapolate age from the earliest reported age. For gender changes, I replace the variable with the most frequently reported gender. Overall, this leaves me with a total of 6,738,077 observations on 654,128 workers.

A.1.2 Consistent Proportional Mapping

I extend SOC 2010 occupational codes over the range of the entire sample period using consistent proportional mapping. First, for each change in occupational coding I create a lookup file, containing the probability for each current SOC to be assigned a particular future SOC. Note that for each current SOC, the probabilities to be assigned to particular subsequent SOCs sum to 1. The lookup file is computed using annual ASHE or NESPD datasets for years in which SOC conversions took place. For the conversion from SOC 2000 to SOC 2010 two datasets exist, each containing occupational codes in either the previous or subsequent SOC coding. Combining these datasets allows computing the probability with which workers with a specific previous SOC are assigned an occupational code corresponding to the subsequent SOC code. For earlier conversions, lookup files are based on workers who stay in the same job in the year before and after the SOC change. Second, I assign each worker a uniform random number between 0 and 1. Random numbers are assigned deterministically such that workers always exhibits the same random number. Third, for each change in occupational coding, based on the assigned uniform random numbers using each individuals unique identifier (piden) I convert current into subsequent SOCs using conversion probabilities contained in the lookup file. In particular, workers are assigned a particular subsequent SOC if their random number lies in the interval corresponding to this conversion's probability. This assures that SOCs are converted in proportion to observed SOC changes. For instance, if in 2010 30 percent of workers with a SOC 2000 code 110 had SOC 2010 code 210, in each year before 2010 30 percent of workers with SOC 2000 code 110 are assigned SOC 2010 code 210. Conversions are consistent as workers are always assigned the same random number, so a worker with SOC 2000 code 110 will consistently be assigned SOC 2010 code 210. To compute a variable containing workers' SOC 2010 codes for all years, I recode SOC 75 codes to SOC 90 codes, original and converted SOC 90 to SOC 2000 codes, and eventually original and converted SOC 2000 to SOC 2010 codes.

The following issues arise during this conversion: In a small number of cases there are observations for SOCs in years before the SOC change, but no observations for these SOCs in the duplicate datasets used to create the lookup file. This is the case for the conversion from SOC 1975 to SOC 1990 for codes 101 and 356, and from SOC 1990 to SOC 2000 for codes 593 and 957. I recode these codes to the closest alternative. The number of observations affected is less than 0.5 percent. Also, for NESPD in some cases codes are reported in previous years which cannot be found in the SOC documentation, and thus most likely constitute typos. I treat these as missing SOC codes and replace them with the last non-missing value.

As a technical note, in the lookup file current SOCs may be assigned to subsequent codes with the same probability. As the analysis was conducted in Stata, which sorts variables with duplicate values in a different order each time sorting is done, the resulting distribution of SOC 2010 codes depends on the order values are sorted within these variables. To get reproducible results, I chose to sort with a stable order. While results for the mapping of occupations into median wage deciles appear to differ slightly each time codes were run, the resulting median wage deciles always exhibited job polarization, i.e. employment for the lowest decile as well as the highest decile was always increasing in relative and absolute terms, and middling deciles generally declined in relative and absolute terms. Note that the order in which Stata sorts these variables when using the stable sorting option, and thus the resulting distribution of median wage deciles, is beyond the control of the researcher.

A.2 Tables

Table A.1: The 1-digit SOC10 Composition of Low, Medium, and High Skill Categories

Low Skilled Employment

SOC 10 Major Group	Percentage
Elementary	29.3
Sales	26.1
Personal Care	23.5
Office/Admin	15.6
Skilled Trade	5.6

Medium Skilled Employment

SOC 10 Major Group	Percentage
Office/Admin	21.2
Skilled Trade	16.0
Associates/Technicians	16.0
Operators/Production	15.6
Elementary	12.8
Managers	6.1
Professionals	5.8
Sales	4.6
Personal Care	2.0

High Skilled Employment

SOC 10 Major Group	Percentage
Professionals	67.3
Managers	22.5
Associates/Technicians	8.4
Operators/Production	1.7

Composition of low, medium and high skilled employment categories over sample period 1975 to 2015 in terms of 1-digit SOC 10 major groups. Workers are assigned to low (medium, high) skilled employment if they are employed in a 3-digit SOC 10 occupation assigned to the 1st (2-8th, 9-10th) decile of the 1975 3-digit occupational median wage distribution. Wage is measured as gross weekly pay. Occupations have been converted to SOC 10 in all years using consistent proportional mapping.

Table A.2: Group Contributions to Decomposition of Low Skilled Employment Share

Age	Period	1975-1979	1979-1981	1981-1990	1990-1992	1992-2008	2008-2009	2009-2015	1981-2015
	Total Change	1.7	0.5	0.8	1.1	2.8	0.4	3.8	8.8
Age	Contribution of female workers								
	Composition	12.9	32.5	90.7	-8.0	-28.3	-24.1	3.9	-0.9
18-30	Propensity	4.2	23.2	3.2	30.0	27.2	38.8	32.5	27.2
	Interaction	0.3	1.0	0.6	-0.6	-4.6	-0.8	1.0	-0.6
	Composition	32.8	32.8	157.9	63.4	34.8	34.7	-9.7	31.7
31-50	Propensity	26.6	19.8	-99.1	13.0	-21.0	-13.9	17.3	-7.9
	Interaction	2.6	0.5	-17.9	1.3	-2.6	-0.2	-0.8	-3.3
	Composition	11.2	-24.8	-54.5	14.9	56.1	50.7	17.8	25.5
51-65	Propensity	9.2	13.6	-8.9	13.9	-16.5	-13.7	13.3	-0.9
	Interaction	0.5	-0.4	1.0	0.7	-7.4	-0.6	1.9	-0.5
Age	Contribution of male workers								
	Composition	-3.0	-2.8	-8.0	-14.2	-15.9	-35.1	2.3	-0.6
18-30	Propensity	12.8	19.3	58.5	10.1	62.6	33.6	13.2	39.9
	Interaction	-0.5	-0.2	-2.3	-0.8	-14.7	-1.6	0.4	-13.3
	Composition	-2.6	6.7	-4.3	0.4	-3.0	-6.3	-4.6	-2.9
31-50	Propensity	1.2	7.9	21.1	-16.9	26.1	18.4	9.7	13.5
	Interaction	0.0	0.1	-0.4	0.0	-1.3	-0.2	-0.7	-1.9
	Composition	-7.9	-19.0	-38.5	-6.6	3.6	4.4	1.1	-2.5
51-65	Propensity	0.4	-11.1	1.2	-0.7	4.4	15.6	1.2	3.4
	Interaction	0.0	0.7	-0.3	0.1	0.5	0.2	0.0	-0.6

Decomposition contributions to change in low skilled employment share in percent. Decomposition is constructed according to equation 3.5. Six groups are considered, three age groups (18-30, 31-50, 51-65) and two gender groups. Results show group's composition, propensity, or interaction term as percentage of respective employment share change. Based on annual observations of NESPD and ASHE. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Sub-periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Table A.3: Group Contributions to Decomposition of High Skilled Employment Share

	Period	1975-1979	1979-1981	1981-1990	1990-1992	1992-2008	2008-2009	2009-2015	1981-2015
	Total Change	0.6	0.5	1.0	0.9	2.0	0.5	-0.4	3.9
Age		Contribution of female workers							
	Composition	16.4	12.1	27.4	-3.4	-13.8	-6.3	-12.8	-0.8
18-30	Propensity	-13.4	-7.2	-12.1	7.8	23.7	13.9	76.1	1.7
	Interaction	-0.9	-0.3	-2.3	0.2	-4.0	-0.3	2.4	0.0
	Composition	29.4	9.8	41.3	33.8	22.6	12.5	45.8	23.5
31-50	Propensity	41.5	13.3	38.8	19.2	11.8	32.3	-83.0	27.2
	Interaction	4.1	0.3	7.0	1.9	1.5	0.5	3.6	11.2
	Composition	6.1	-4.4	-8.8	4.7	20.3	13.4	-61.5	11.7
51-65	Propensity	9.9	10.6	18.1	1.7	14.7	10.8	67.5	8.6
	Interaction	0.5	-0.3	-2.1	0.1	6.6	0.5	9.8	5.2
Age		Contribution of male workers							
	Composition	-18.0	-4.5	-10.6	-21.5	-27.8	-15.0	-12.5	-22.7
18-30	Propensity	-11.5	6.3	-12.9	7.7	-2.3	16.1	69.1	-11.2
	Interaction	0.4	-0.1	0.5	-0.6	0.5	-0.8	2.2	3.8
	Composition	-21.1	18.0	-10.2	1.5	-15.8	-11.9	107.8	-19.2
31-50	Propensity	51.5	42.7	59.3	52.9	22.9	24.9	-66.4	50.1
	Interaction	-1.3	0.8	-1.1	0.1	-1.1	-0.3	4.8	-7.0
	Composition	-32.9	-27.9	-54.2	-15.9	10.8	7.1	-22.2	-9.9
51-65	Propensity	42.7	33.0	28.5	11.2	26.5	2.5	-29.7	33.5
	Interaction	-3.6	-2.1	-6.6	-0.8	2.9	0.0	-1.0	-5.7

Decomposition contributions to change in high skilled employment share in percent. Decomposition is constructed according to equation 3.5. Six groups are considered, three age groups (18-30, 31-50, 51-65) and two gender groups. Results show group's composition, propensity, or interaction term as percentage of respective employment share change. Based on annual observations of NESPD and ASHE. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Sub-periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Table A.4: Comparing Benchmark and Actual Employment Share Changes

		1981	1981-1990		1981-2008		1981-2015	
		Level	Change					
		Actual	Actual	Average	Actual	Average	Actual	Average
Age	Status	Female workers						
18-30	Low	28.5	0.2	0.4	7.2	6.9	18.1	18.7
	Medium	60.5	0.7	0.9	-9.8	-8.4	-18.6	-19.0
	High	11.0	-0.9	-1.3	1.5	1.5	0.5	0.3
31-50	Low	37.6	-4.3	-2.8	-6.1	-5.5	-3.9	-3.5
	Medium	49.8	2.1	0.7	2.2	2.2	-2.1	-2.4
	High	12.5	2.1	2.0	3.9	3.3	6.0	6.0
51-65	Low	42.5	-0.8	-1.0	-4.6	-3.6	-0.9	-1.8
	Medium	48.8	-1.3	-1.0	-1.3	-1.6	-3.0	-2.5
	High	8.8	2.1	2.0	5.8	5.2	3.9	4.3
Age	Status	Male workers						
18-30	Low	8.7	2.5	2.3	14.0	10.8	19.3	20.4
	Medium	76.5	-1.8	-1.1	-13.4	-10.1	-16.8	-17.9
	High	14.7	-0.7	-1.2	-0.6	-0.7	-2.4	-2.5
31-50	Low	6.7	0.6	0.7	2.7	1.8	4.4	4.4
	Medium	73.3	-2.8	-2.3	-8.3	-6.2	-11.8	-11.5
	High	20.0	2.2	1.6	5.6	4.4	7.3	7.0
51-65	Low	8.5	0.1	-0.2	1.1	0.7	2.0	1.8
	Medium	76.3	-1.9	-1.6	-8.7	-7.0	-10.8	-11.0
	High	15.2	1.9	1.8	7.6	6.2	8.8	9.2

Comparison of actual and benchmark employment share changes. Actual employment share changes are computed based on equation 2.2, benchmark employment share changes are computed based on equation 2.3. Six groups are considered: three age groups (18-30, 31-50, 51-65), and two gender groups. Actual employment share changes are based on annual transition rates, benchmark changes on average period transition rates. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table A.5: Counterfactual Results for Prime Aged Female Workers

Period	1981-1990				1981-2008				1981-2015			
					Low Skilled Employment							
BM change:	-2.8				-5.5				-3.5			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-4.1	-1.0	0.0	-2.8	-4.0	-0.5	0.1	-3.1	-2.5	-0.3	0.0	-2.3
Med	1.8	4.2	0.0	2.7	1.5	5.5	0.0	3.9	0.7	2.9	0.0	2.4
High	0.1	0.0	1.0	0.8	0.2	0.0	1.0	0.9	0.3	0.0	0.7	0.4
Non	10.8	-8.1	-2.1	0.1	17.0	-11.1	-3.5	0.4	12.8	-9.3	-2.0	0.2
Non (cr)	2.7	-2.2	-0.2	0.0	7.2	-5.1	-0.8	0.2	3.9	-3.5	0.4	0.0
Entry	n.a.				0.9				0.2			
					Medium Skilled Employment							
BM change:	0.7				2.2				-2.4			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-3.2	-1.0	0.0	-2.1	-3.1	-0.5	0.0	-2.3	1.8	0.3	0.0	1.6
Med	1.8	6.0	0.2	3.7	1.5	8.3	-0.3	5.8	-0.7	-4.6	0.6	-3.7
High	0.0	-0.6	-1.6	-1.1	0.0	-0.6	-1.7	-1.3	0.0	0.5	1.0	0.5
Non	8.2	-11.4	2.9	0.0	12.6	-16.3	5.1	0.2	-8.8	14.3	-2.4	-0.1
Non (cr)	2.1	-3.0	0.2	0.0	5.4	-7.5	1.1	0.1	-2.7	5.3	0.5	0.0
Entry	n.a.				0.3				1			
					High Skilled Employment							
BM change:	2.0				3.3				6.0			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-0.9	0.0	0.0	-0.7	-0.9	0.0	0.1	-0.8	-0.7	0.0	0.0	-0.7
Med	0.0	-1.8	-0.2	-1.1	0.0	-2.8	0.3	-1.8	0.0	-1.6	0.6	-1.3
High	0.1	0.6	2.7	2.0	0.2	0.6	2.7	2.1	0.3	0.5	1.6	0.9
Non	2.6	3.3	-5.0	0.1	4.4	5.2	-8.6	0.1	4.1	5.0	-4.4	0.1
Non (cr)	0.6	0.9	-0.4	0.0	1.9	2.4	-1.9	0.1	1.3	1.9	0.9	0.0
Entry	n.a.				0.7				1.1			

Contributions to employment share changes for prime aged female workers based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A, B \in$ (Low, Med, High, Non, Entry). Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Prime aged female workers comprise workers aged 31-50. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table A.6: Counterfactual Results for Older Female Workers

Period	1981-1990				1981-2008				1981-2015			
					Low Skilled Employment							
BM change:	-1.0				-3.6				-1.8			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-3.4	-1.6	-0.1	-2.1	-4.8	-0.5	-0.1	-4.3	-4.1	0.4	-0.1	-4.8
Med	1.8	2.8	0.0	1.8	0.7	4.5	0.0	3.9	-0.4	3.8	0.0	4.1
High	0.3	0.0	0.9	0.6	0.5	0.1	1.7	1.3	0.4	0.0	1.1	0.9
Non	19.0	-15.6	-4.0	0.8	23.0	-16.7	-4.8	1.2	18.3	-14.6	-3.3	1.1
Non (cr)	4.5	-4.3	-0.7	0.3	7.0	-5.8	-0.5	0.6	3.7	-4.0	0.1	0.4
Entry	n.a.				1.1				0.4			
					Medium Skilled Employment							
BM change:	-1.0				-1.6				-2.5			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	2.8	1.6	0.0	1.7	3.7	0.5	0.0	3.2	3.3	-0.4	0.0	3.6
Med	-1.8	-3.6	-0.6	-2.2	-0.7	-6.1	-0.3	-5.3	0.4	-5.0	0.0	-5.4
High	0.0	0.7	1.1	0.7	0.0	1.0	2.0	1.4	0.0	0.6	1.3	1.0
Non	-15.5	19.5	-4.4	-0.1	-17.4	22.7	-5.3	-0.4	-13.8	19.3	-3.3	-0.9
Non (cr)	-3.7	5.3	-0.7	-0.1	-5.3	7.9	-0.6	-0.2	-2.8	5.3	0.1	-0.3
Entry	n.a.				0.5				0.4			
					High Skilled Employment							
BM change:	2.0				5.2				4.3			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-0.6	0.0	-0.1	-0.4	-1.1	0.0	-0.1	-1.0	-0.9	0.0	-0.1	-1.2
Med	0.0	-0.7	-0.6	-0.4	0.0	-1.6	-0.4	-1.4	0.0	-1.2	0.0	-1.3
High	0.3	0.7	2.0	1.3	0.5	1.0	3.7	2.7	0.4	0.6	2.4	1.9
Non	3.5	3.9	-8.4	0.6	5.7	6.0	-10.1	0.8	4.5	4.7	-6.6	0.3
Non (cr)	0.8	1.1	-1.4	0.3	1.7	2.1	-1.1	0.4	0.9	1.3	0.2	0.1
Entry	n.a.				1.6				0.8			

Contributions to employment share changes for older female workers based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A, B \in$ (Low, Med, High, Non, Entry). Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Older female workers comprise workers aged 51-65. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table A.7: Counterfactual Results for Medium Skilled Female Workers

Period	1981-1990				1981-2008				1981-2015			
					Young workers							
BM change:	0.9				-8.4				-19.0			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Non low	2.8	-1.0	0.0	0.7	-3.1	0.8	0.0	-0.4	-0.7	0.6	0.4	0.0
Non med	0.6	-3.9	0.3	-0.4	-0.5	2.8	0.1	-0.3	0.7	1.5	0.5	-0.6
Non high	0.1	0.1	0.8	0.2	-0.4	-0.4	-0.9	-0.3	-0.1	-0.1	-0.1	0.1
Non low (cr)	0.4	0.2	0.0	0.2	-0.4	-1.1	0.1	-0.3	1.7	-1.4	0.5	0.2
Non med (cr)	-0.3	-0.1	-0.1	0.1	0.2	-1.0	0.5	-1.2	1.1	-1.8	0.7	-1.6
Non high (cr)	0.1	0.2	0.2	0.1	-0.3	-0.6	-0.3	-0.2	0.0	-0.3	0.3	0.1
					Prime aged workers							
BM change:	0.8				2.0				-2.4			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Non low	5.3	-1.5	0.0	1.4	7.8	-3.4	0.1	1.3	-5.5	3.1	0.0	-0.9
Non med	1.8	-8.1	0.3	-2.1	3.6	-11.1	0.5	-1.5	-2.6	9.2	0.1	0.9
Non high	0.1	-0.6	2.1	0.8	0.3	-1.2	3.5	0.8	-0.2	1.1	-2.8	-0.6
Non low (cr)	1.5	0.0	-0.1	0.5	3.6	-1.1	0.0	0.8	-1.8	1.0	0.0	-0.2
Non med (cr)	0.3	-2.5	-0.1	-0.8	1.4	-5.3	-0.2	-0.9	-0.8	3.3	0.7	0.0
Non high (cr)	0.0	0.0	0.6	0.3	0.1	-0.2	1.4	0.3	-0.1	0.2	-0.8	-0.1
					Older workers							
BM change:	-1.1				-1.6				-2.5			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
N low	-11.5	3.7	-0.2	-6.3	-13.1	4.4	-0.3	-4.8	-10.0	3.5	-0.1	-3.9
Non med	-3.9	15.1	-1.1	6.3	-4.6	17.5	-1.5	5.4	-3.4	14.8	-0.9	3.6
Non high	-0.4	1.3	-2.9	-1.2	-0.6	2.5	-4.9	-2.0	-0.5	2.0	-3.9	-1.4
Non low (cr)	-3.0	0.5	-0.1	-2.4	-4.3	0.5	-0.2	-2.2	-2.2	0.2	0.0	-1.0
Non med (cr)	-0.5	4.2	-0.1	2.3	-0.6	6.0	-0.1	2.4	0.0	4.1	0.1	0.9
Non high (cr)	-0.1	0.2	-0.7	-0.4	-0.1	0.7	-1.2	-0.8	-0.1	0.5	-0.9	-0.3

Contributions to medium skilled employment share changes for female workers based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A \in$ (Non low, Non med, Non high) and $B \in$ (Low, Med, High, Non). Non low (med, high) refers to non-employed workers whose previous job was low (med, high) skilled. Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Three age groups are considered: Young (18-30 years), prime aged (31-50), older (51-65). Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table A.8: Counterfactual Results for Medium Skilled Male workers

Period	1981-1990				1981-2008				1981-2015			
					Young workers							
BM change:	-1.2				-10.5				-17.9			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Non low	-1.1	0.6	-0.2	-0.3	-2.4	1.9	-0.4	-0.2	-1.3	2.0	-0.3	-0.1
Non med	-0.5	3.6	-1.3	0.7	-0.2	5.1	-0.8	0.9	0.4	4.7	-0.4	1.0
Non high	0.0	0.2	-1.1	-0.3	0.0	0.4	-1.1	-0.2	0.0	0.3	-0.6	-0.1
Non low (cr)	-0.2	0.2	0.0	-0.1	-0.7	0.7	-0.1	-0.1	0.3	0.6	0.0	0.0
Non med (cr)	0.1	1.1	-0.1	0.3	0.3	2.0	0.2	0.2	0.9	1.4	0.3	0.0
Non high (cr)	0.0	0.0	-0.3	-0.1	0.0	0.1	-0.4	-0.1	0.0	0.1	-0.1	0.1
					Prime aged workers							
BM change:	-2.4				-6.5				-11.4			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Non low	-1.5	0.8	-0.2	-0.4	-3.5	2.0	-1.0	-0.4	-3.1	2.4	-0.9	-0.6
Non med	-1.0	6.5	-2.2	1.9	-2.0	10.9	-4.6	2.1	-1.3	11.7	-2.8	1.8
Non high	-0.1	0.8	-3.9	-1.3	-0.2	2.4	-6.4	-1.7	-0.2	2.5	-5.8	-1.8
Non low (cr)	-0.4	0.1	0.0	-0.1	-1.4	0.9	-0.4	-0.3	-1.0	1.1	-0.4	-0.3
Non med (cr)	0.0	2.2	0.1	0.7	-0.5	6.0	-0.5	1.3	-0.2	6.0	0.7	0.7
Non high (cr)	0.0	-0.1	-0.9	-0.3	-0.1	0.7	-2.0	-0.6	-0.1	0.8	-1.5	-0.5
					Older workers							
BM change:	-1.7				-7.2				-11.0			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Non low	-3.4	0.9	-0.2	-1.5	-5.4	1.8	-0.5	-1.6	-4.7	2.0	-0.4	-1.7
Non med	-2.6	9.1	-2.7	3.4	-3.8	14.5	-4.4	5.0	-3.1	15.3	-2.9	4.7
Non high	-0.3	1.4	-6.4	-2.6	-0.7	3.7	-12.2	-4.3	-0.6	3.9	-10.6	-4.1
Non low (cr)	-0.9	0.2	0.0	-0.7	-1.8	0.6	-0.1	-1.0	-1.3	0.7	0.0	-1.0
Non med (cr)	-0.4	3.3	-0.3	1.7	-0.8	7.2	-0.8	3.5	-0.7	7.2	0.2	3.1
Non high (cr)	-0.1	0.3	-1.9	-1.1	-0.3	1.5	-4.5	-2.7	-0.3	1.6	-3.3	-2.1

Contributions to medium skilled employment share changes for male workers based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A \in$ (Non low, Non med, Non high) and $B \in$ (Low, Med, High, Non). Non low (med, high) refers to non-employed workers whose previous job was low (med, high) skilled. Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Three age groups are considered: Young (18-30 years), prime aged (31-50), older (51-65). Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table A.9: Aggregated Counterfactual Group Results

Period	1981-1990				1981-2008				1981-2015			
	Low skilled employment											
BM change:	1.0				4.2				8.9			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	2.4	1.0	0.1	1.3	2.4	0.9	0.1	1.4	1.8	0.7	0.2	1.0
Med	-1.1	-2.4	0.0	-1.6	-0.8	-2.4	0.0	-2.0	0.1	-0.9	0.0	-1.4
High	-0.1	0.0	-0.6	-0.4	-0.2	0.0	-0.7	-0.6	-0.3	0.0	-0.5	-0.4
Non	-6.3	4.9	1.4	-0.2	-10.6	7.5	2.3	-0.2	-7.6	7.2	1.7	-0.2
Non (cr)	-1.1	1.2	0.2	-0.1	-3.2	3.0	0.4	-0.3	-0.8	2.4	-0.2	-0.3
Entry	0.6	0.2	0.1	0.0	1.6	0.9	-0.1	-0.3	3.4	2.4	0.7	0.4
	Medium skilled employment											
BM change:	-1.7				-7.2				-12.9			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	2.0	1.0	0.0	1.1	1.9	0.9	0.0	1.0	1.3	0.7	0.0	0.7
Med	-1.2	-4.8	-0.7	-3.1	-0.8	-4.4	-0.4	-3.5	0.1	-1.6	0.1	-2.2
High	0.0	1.0	2.5	1.7	0.0	1.6	2.5	1.7	0.0	1.3	1.3	0.7
Non	-5.0	9.1	-5.2	0.0	-7.9	14.2	-7.8	0.0	-5.3	13.8	-5.1	-0.1
Non (cr)	-0.8	2.5	-0.8	0.0	-2.3	6.3	-1.8	0.0	-0.4	5.3	-0.2	-0.1
Entry	0.5	0.3	-0.3	0.0	2.0	1.9	0.2	0.4	3.7	3.6	1.0	1.3
	High skilled employment											
BM change:	0.7				3.0				3.9			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-0.5	0.0	-0.1	-0.3	-0.6	0.0	0.1	-0.4	-0.4	0.0	-0.2	-0.3
Med	0.0	-2.4	-0.8	-1.4	0.0	-2.0	-0.4	-1.4	0.0	-0.7	0.1	-0.8
High	0.1	1.1	3.1	2.2	0.3	1.6	3.2	2.3	0.3	1.4	1.8	1.1
Non	1.3	4.1	-6.6	0.2	2.7	6.8	10.1	0.2	2.3	6.6	-6.8	0.1
Non (cr)	0.2	1.2	-1.0	0.1	0.8	3.3	-2.1	0.2	0.3	3.0	0.0	0.1
Entry	-0.1	0.1	-0.4	0.0	0.4	1.0	0.3	0.7	0.3	1.2	0.3	0.8

Aggregated contributions to employment share changes based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A, B \in$ (Low, Med, High, Non, Entry). Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Six groups are considered, three aged groups (18-30, 31-50, 51-65), and two gender groups. Contributions are aggregated using annual population weights of each group. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

Table A.10: Results for Aggregate Counterfactual Analysis

Period	1981-1990				1981-2008				1981-2015			
	Low skilled employment											
BM change:	1.0				4.2				8.9			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	2.7	1.2	0.2	1.8	2.8	1.4	0.4	1.9	1.6	0.7	0.2	1.0
Med	-0.9	-2.6	0.2	-1.9	-0.1	-2.2	0.3	-1.8	0.6	-1.0	0.0	-1.2
High	0.0	0.1	-0.7	-0.5	0.1	0.2	-0.9	-0.5	-0.3	-0.1	-0.9	-0.5
Non	-7.3	6.4	2.2	0.3	-10.4	9.8	4.1	1.3	-7.4	10.5	3.4	1.3
Non (cr)	-0.2	1.3	0.3	0.2	-1.4	4.1	0.8	0.7	0.6	3.9	-0.1	0.5
Entry	0.9	0.7	0.3	0.1	2.8	2.1	0.5	0.2	4.2	3.0	0.4	0.0
	Medium skilled employment											
BM change:	-1.7				-7.2				-12.9			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	2.2	1.2	0.1	1.5	2.0	1.2	0.1	1.3	1.2	0.6	0.0	0.7
Med	-0.9	-4.9	-0.7	-3.5	-0.3	-4.5	-0.5	-3.5	0.6	-1.7	-0.1	-2.0
High	0.2	1.4	2.9	2.1	0.1	2.1	3.2	2.0	0.0	1.7	1.9	0.8
Non	-5.8	11.0	-6.2	0.1	-7.8	16.9	-9.0	0.6	-5.1	17.5	-6.0	0.5
Non (cr)	-0.1	2.2	-0.4	0.1	-1.2	6.8	-1.1	0.3	0.4	6.5	0.1	0.2
Entry	0.8	1.0	-0.3	0.1	2.1	2.5	-0.3	0.1	3.3	3.6	-0.4	0.0
	High skilled employment											
BM change:	0.7				3.0				3.9			
TR from/to:	low	med	high	non	low	med	high	non	low	med	high	non
Low	-0.5	0.0	-0.1	-0.3	-0.8	-0.2	-0.3	-0.6	-0.4	0.0	-0.2	-0.3
Med	0.0	-2.3	-0.9	-1.5	-0.2	-2.3	-0.8	-1.8	0.0	-0.7	-0.1	-0.8
High	0.1	1.3	3.6	2.6	0.1	1.9	4.1	2.5	0.3	1.8	2.8	1.3
Non	1.5	4.6	-8.4	-0.1	2.6	7.1	-13.0	-0.8	2.2	7.1	0.2	-0.8
Non (cr)	0.1	0.9	-0.7	0.0	0.2	2.7	-1.9	-0.4	-0.2	2.6	0.2	-0.3
Entry	-0.1	0.3	-0.6	0.0	-0.7	0.5	-0.8	-0.2	-0.9	0.7	-0.8	0.0

Contributions to employment share changes based on counterfactual analysis. Contribution gives mitigation of counterfactual to benchmark employment share change in percentage points. Benchmark employment share changes are computed based on equation 2.3. Counterfactual employment share changes are computed based on equation 2.4. Contributions correspond to some transition rate held constant during expansionary periods at its 1975-1978 average level. Transition rates give probability to change from labour market state A in year t to labour market state B in year $t+1$. State A is given in column 1, state B in columns 2-13, for $A, B \in$ (Low, Med, High, Non, Entry). Suffix (cr) indicates transition rates corrected for sample start problems using spell duration composition for years 1981 to 1989. Annual employment shares are computed as the fraction of employed workers in the respective job type of total employment. Periods starting at through of 1981 recession and last until peak of subsequent recessions in 1990, 2008, or end of sample period. Based on annual observations of NESPD and ASHE.

A.3 Figures

Figure A.1: Wage Distribution of Occupational Median Wage Deciles, 1975 to 2015

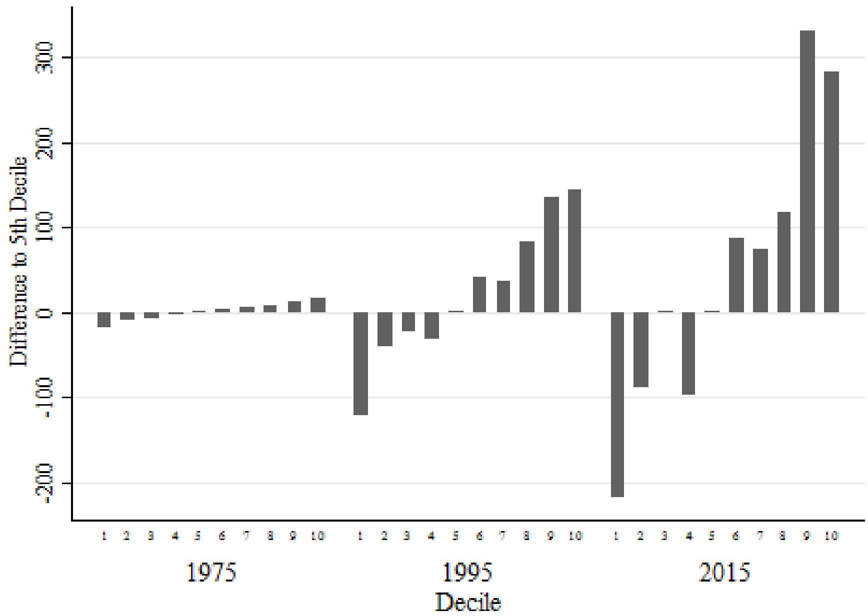

Deviation of 3-digit SOC10 occupational median wage deciles from median wage of 5th decile. Comparisons are made within each year: 1975, 1995, 2015. Figure demonstrates that occupations assigned to low (1st decile), medium (2nd to 8th decile), or high skilled (9th to 10th decile) employment types are in the lowest, middling, or highest median wage groups throughout sample period. Based on annual observations of NESPD and ASHE.

Figure A.2: Occupational Median Wage Decile Employment Growth, 1975 to 2015

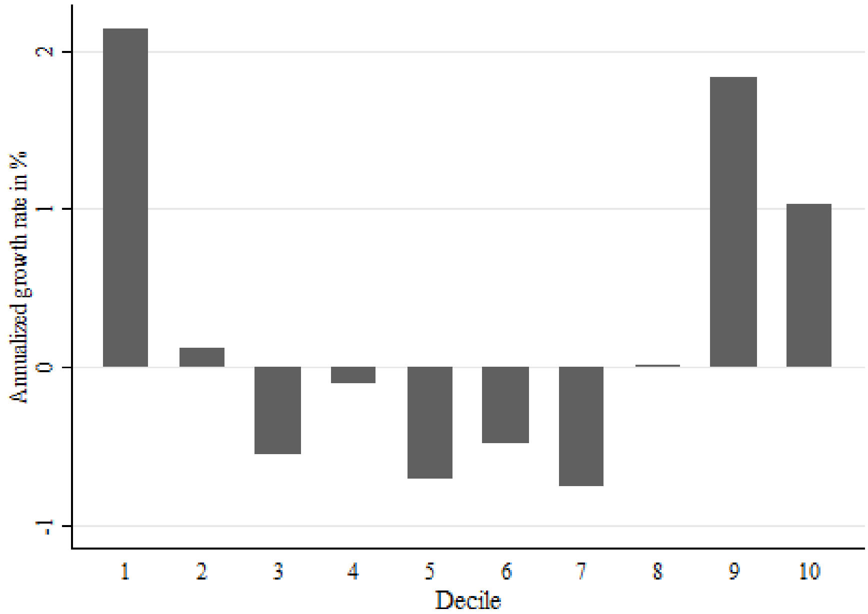

Annualized employment growth rate of median wage deciles over the period 1975 to 2015. Annualized growth rate is computed as percentage change of employment stocks divided by number of years for period 1975 to 2015. Deciles are based on median wage of SOC 10 occupations at 3-digit level in 1975. Stock of employment for each decile corresponds to number of workers in occupations assigned to respecitve decile of the 1975 3-digit SOC 10 occupational median wage distribution. Based on annual observations of NESPD and ASHE.

Figure A.3: Annual Population Weights, 1975 to 2015

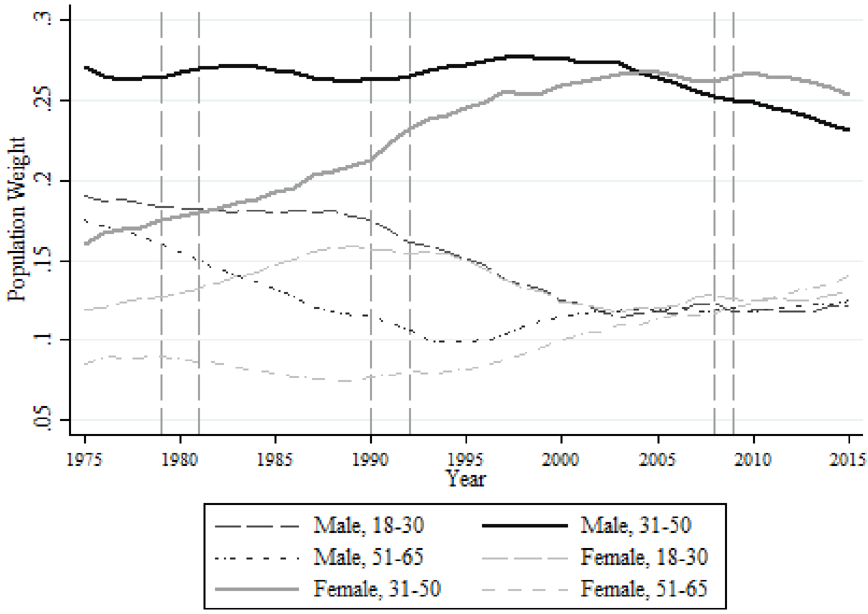

Annual population weights for demographic groups. Six groups are considered: three age groups (18-30, 31-50, 51-65) and two gender groups. Annual population weights are computed as the ratio of the sum of employed workers of the respective group to total employment. Based on annual observations of NESPD and ASHE. Vertical lines indicate recessions.

Figure A.4: Transition rates from low skilled employment for male workers

(a) Low to Low

(b) Low to Med

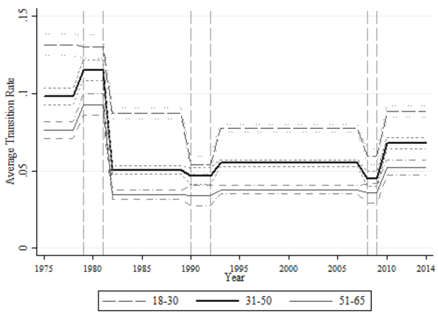

(c) Low to High

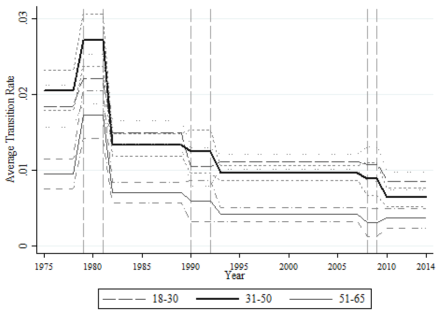

(d) Low to Non

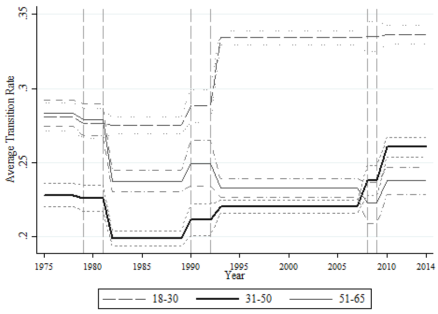

Period average transition rates from low skilled employment for male workers. Transition rates from low to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.5: Transition rates from medium skilled employment for male workers

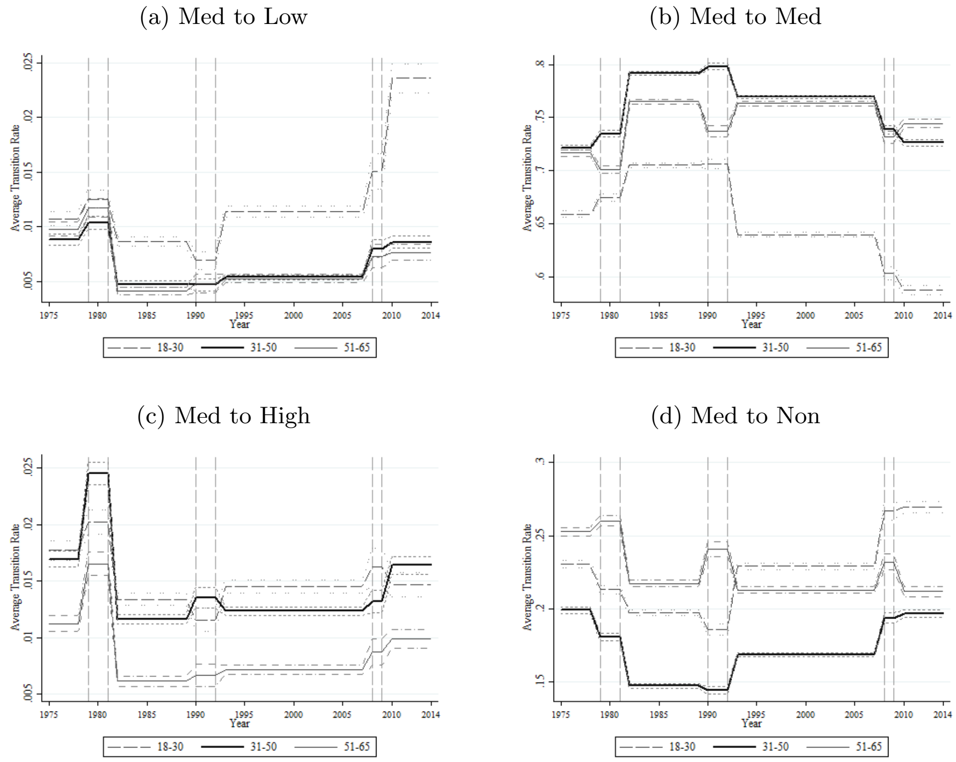

Period average transition rates from medium skilled employment for male workers. Transition rates from medium to low , medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.6: Transition rates from high skilled employment for male workers

(a) High to Low

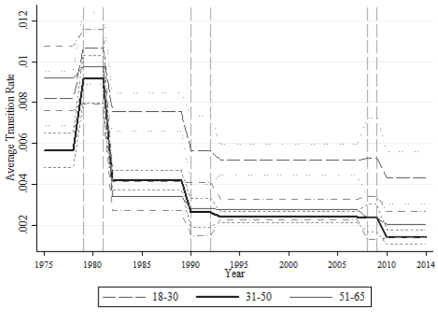

(b) High to Med

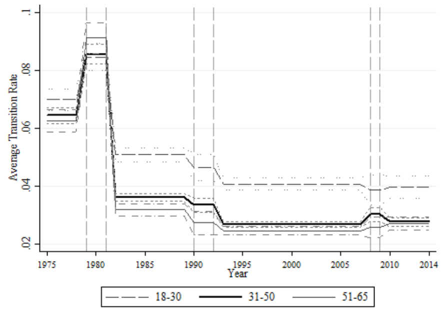

(c) High to High

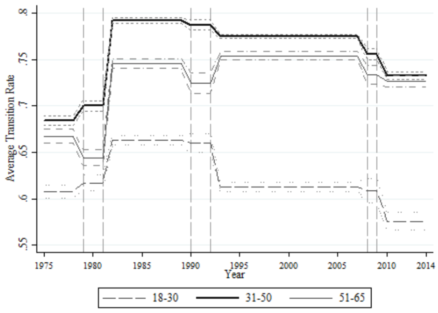

(d) High to Non

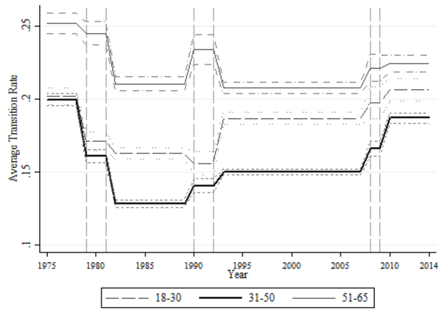

Period average transition rates from high skilled employment for male workers. Transition rates from high to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.7: Transition rates from non-employment for male workers

(a) Non to Low

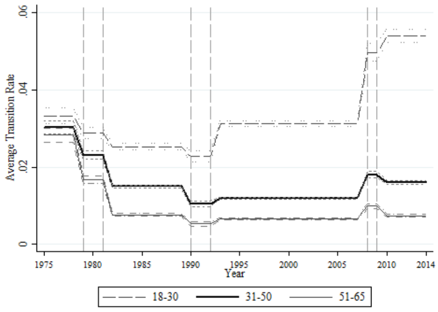

(b) Non to Med

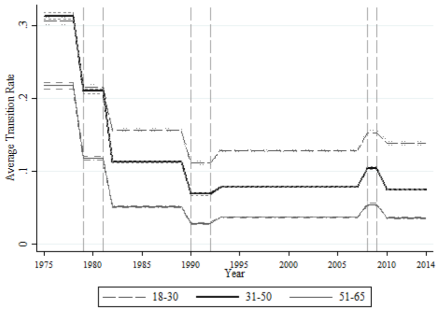

(c) Non to High

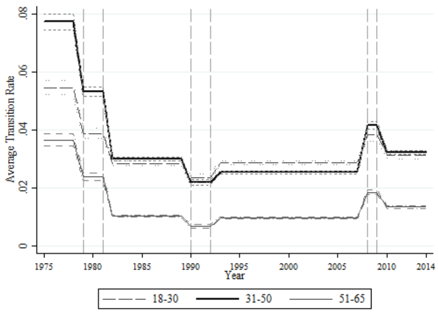

(d) Non to Non

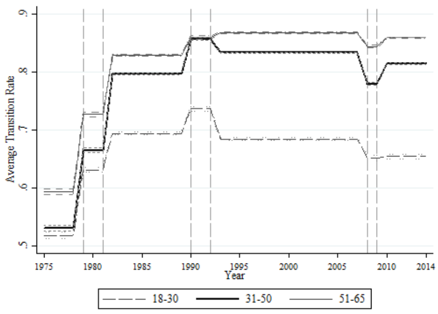

Period average transition rates from non-employment for male workers. Transition rates from non-employment to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.8: Transition rates from sample entry for male workers

(a) Entry to Low

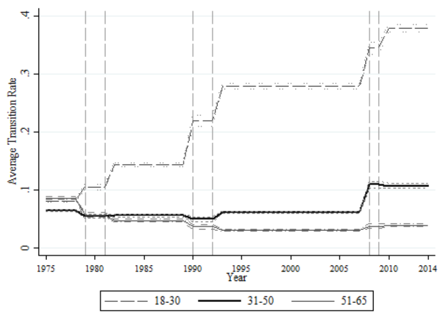

(b) Entry to Med

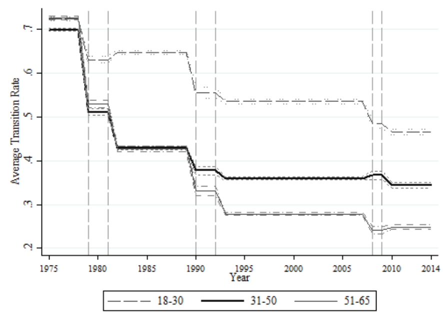

(c) Entry to High

(d) Entry to Non

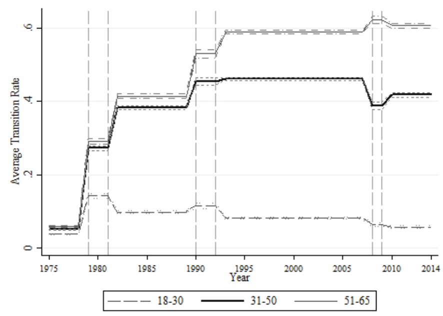

Period average transition rates from sample entry for male workers. Transition rates from entry to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.9: Transition rates from low skilled employment for female workers

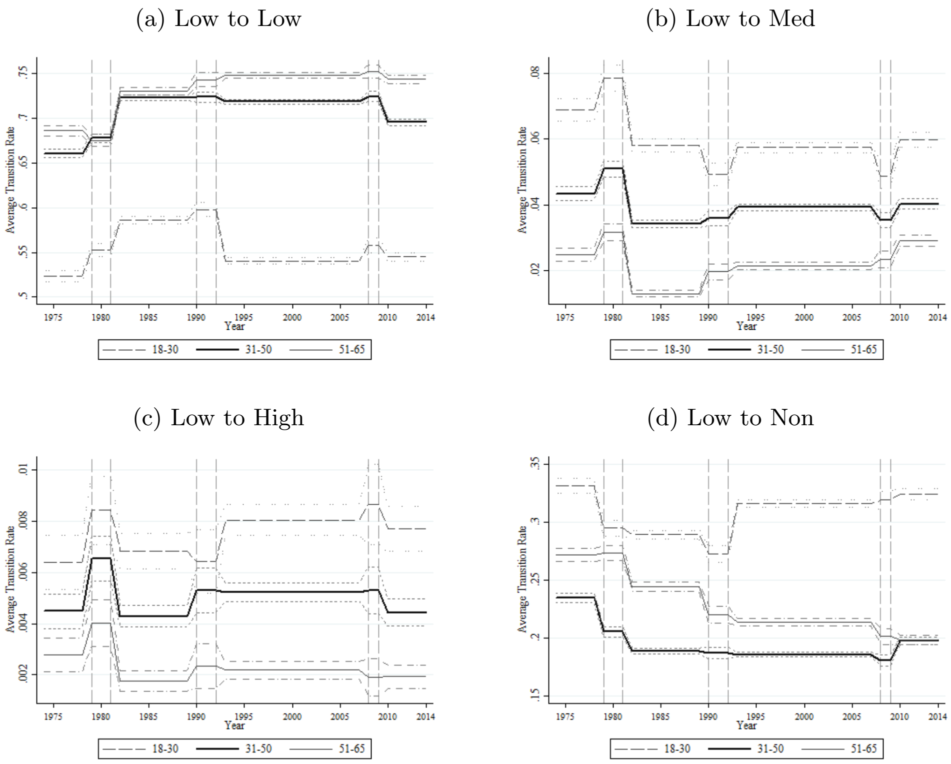

Period average transition rates from low skilled employment for female workers. Transition rates from low to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.10: Transition rates from medium skilled employment for female workers

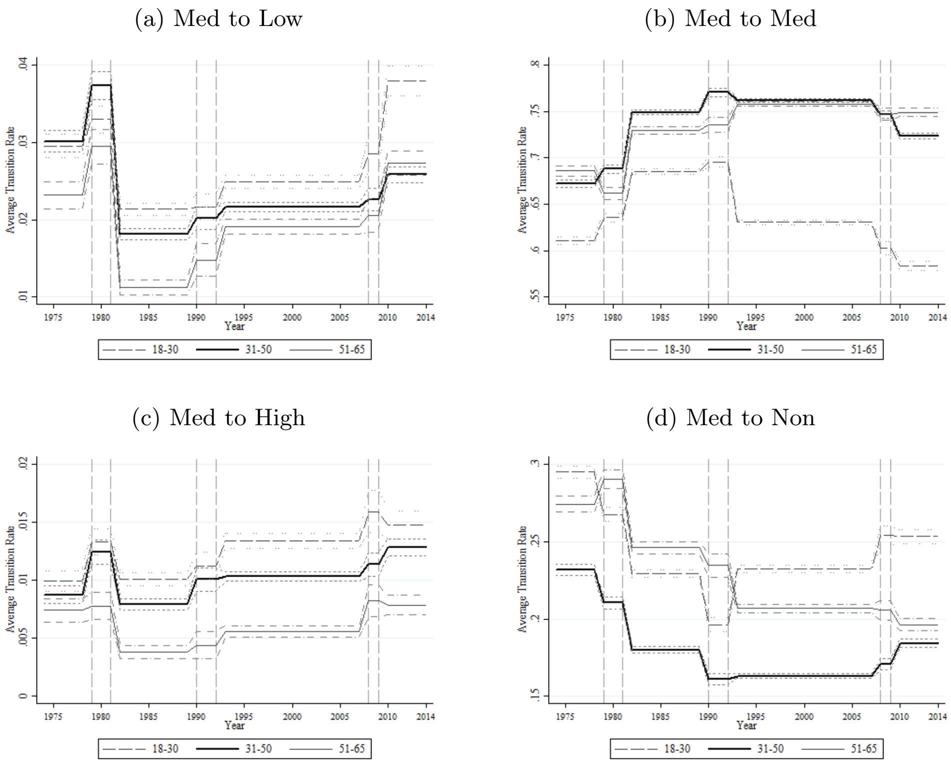

(a) Med to Low

(b) Med to Med

(c) Med to High

(d) Med to Non

Period average transition rates from medium skilled employment for female workers. Transition rates from medium to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.11: Transition rates from high skilled employment for female workers

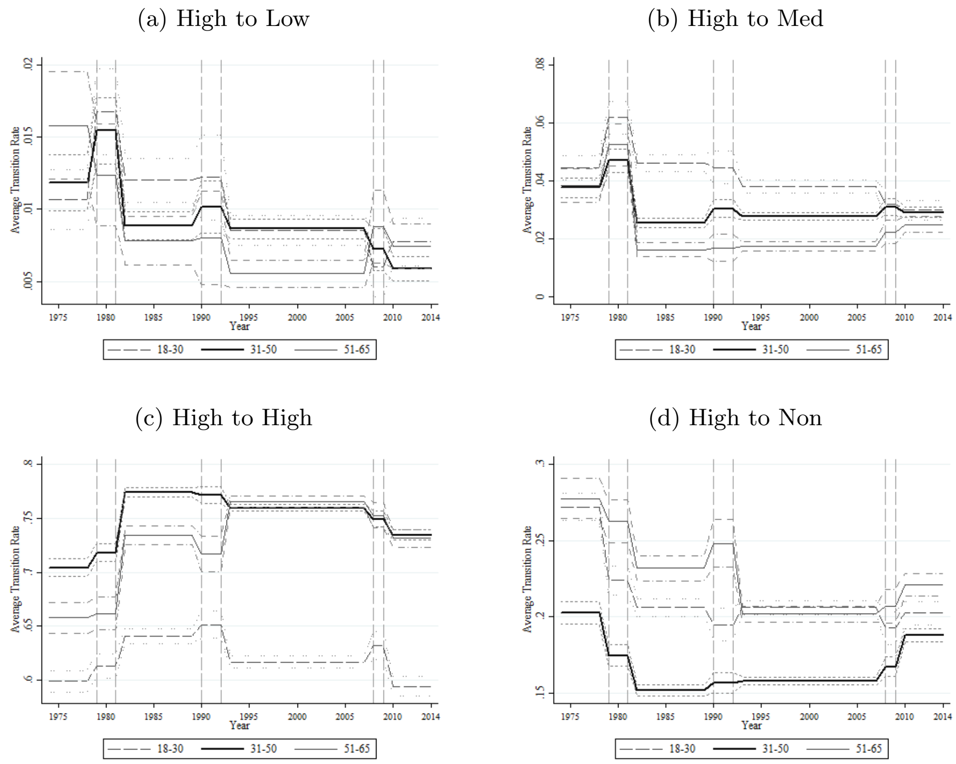

Period average transition rates from high skilled employment for female workers. Transition rates from high to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.12: Transition rates from non-employment

(a) Non to Low

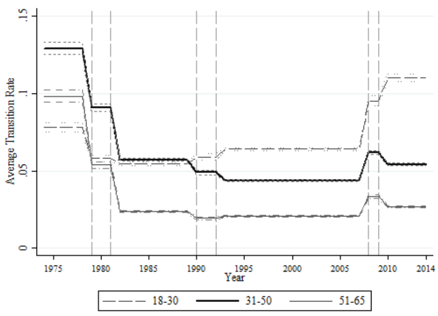

(b) Non to Med

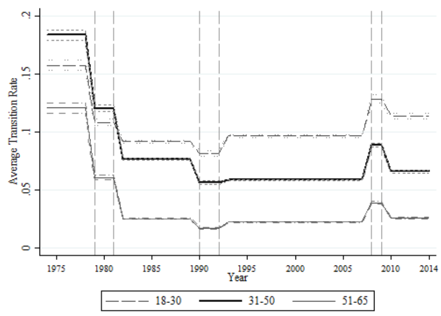

(c) Non to High

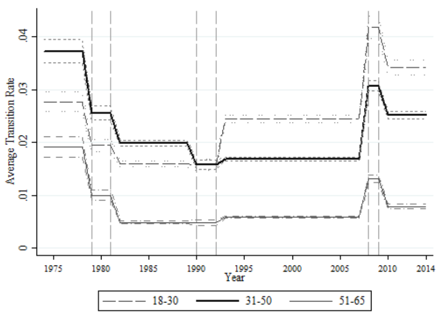

(d) Non to Non

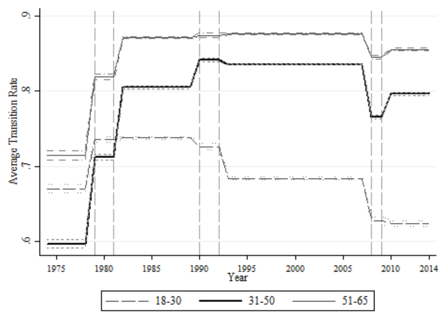

Period average transition rates from non-employment for female workers. Transition rates from non-employment to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.13: Transition rates from sample entry for male workers

(a) Entry to Low

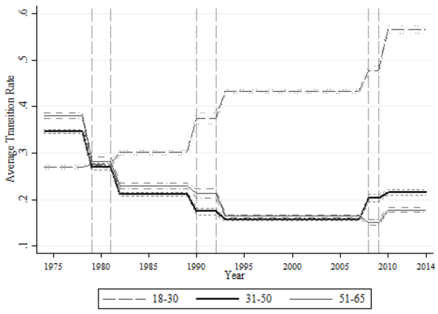

(b) Entry to Med

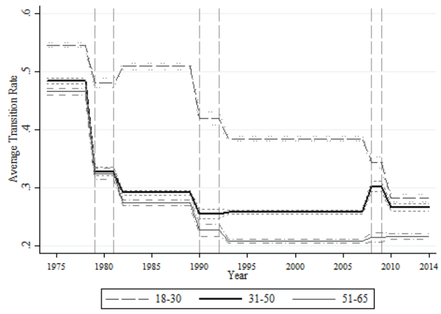

(c) Entry to High

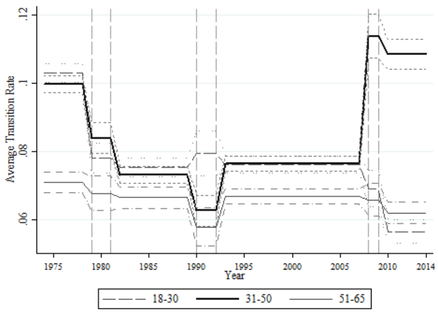

(d) Entry to Non

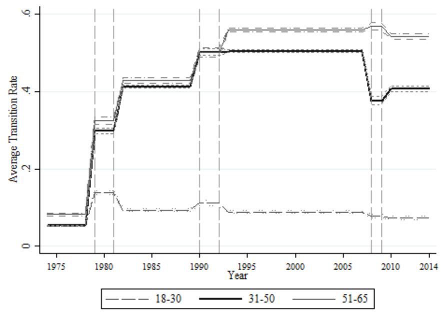

Period average transition rates from sample entry for female workers. Transition rates from entry to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.14: Non-employment outflow rates for male workers previously in low skilled job

(a) Non-Low to Low

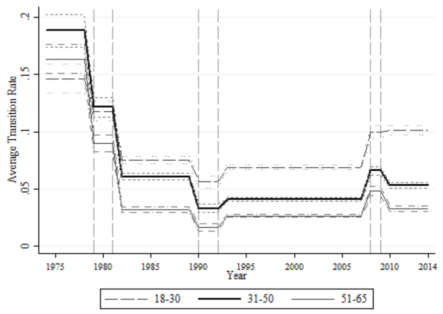

(b) Non-Low to Med

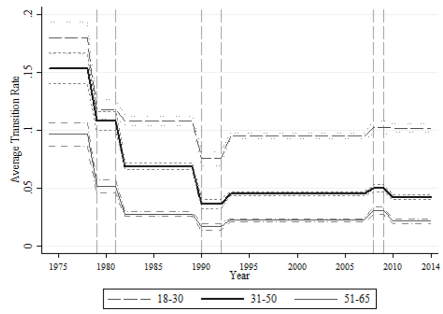

(c) Non-Low to High

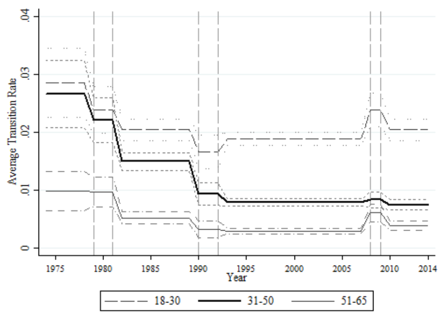

(d) Non-Low to Non-Low

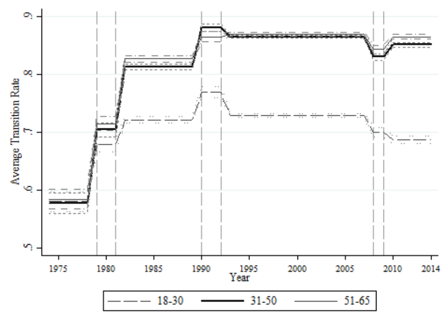

Period average transition rates from non-employment for male workers entering non-employment from low skilled job. Transition rates from low skilled non-employment to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.15: Non-employment outflow rates for male workers previously in medium skilled job

(a) Non-Med to Low

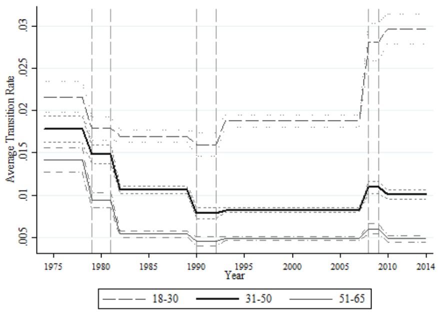

(b) Non-Med to Med

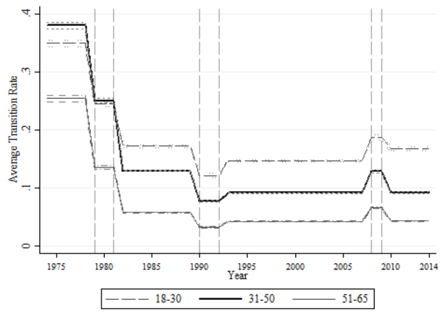

(c) Non-Med to High

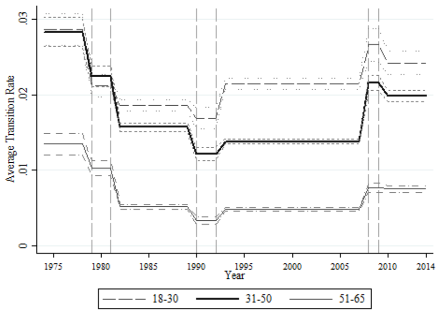

(d) Non-Med to Non-Med

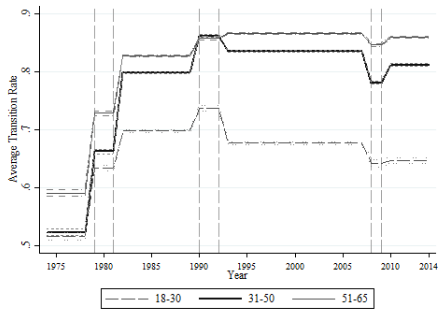

Period average transition rates from non-employment for male workers entering non-employment from medium skilled job. Transition rates from medium skilled non-employment to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.16: Non-employment outflow rates for male workers previously in high skilled job

(a) Non-High to Low

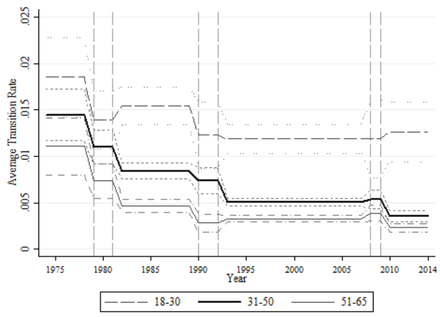

(b) Non-High to Med

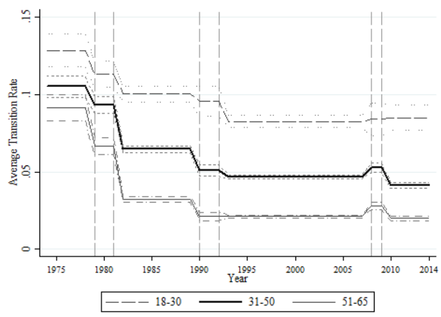

(c) Non-High to High

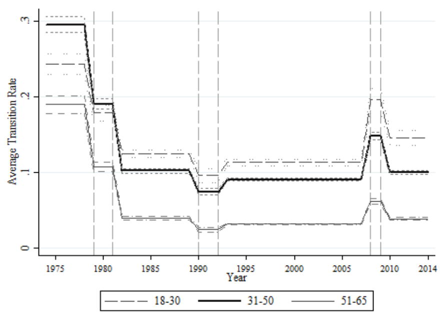

(d) Non-High to Non-High

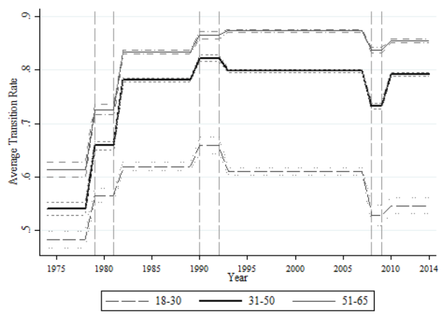

Period average transition rates from non-employment for male workers entering non-employment from high skilled job. Transition rates from high skilled non-employment to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.17: Non-employment outflow rates for female workers previously in low skilled job

(a) Non-Low to Low

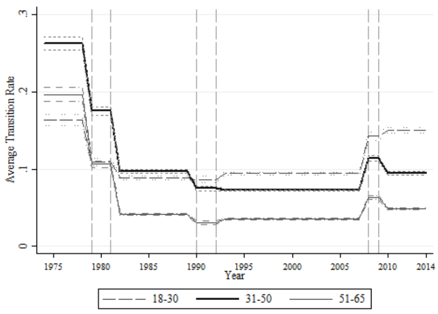

(b) Non-Low to Med

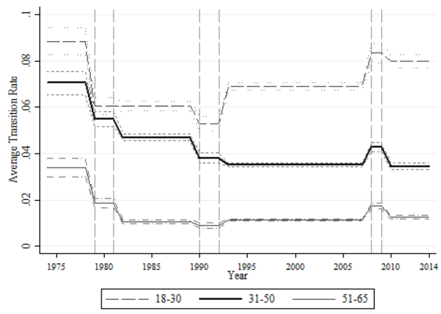

(c) Non-Low to High

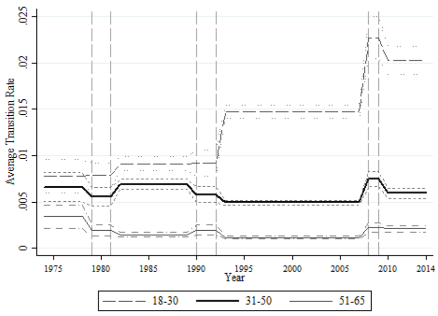

(d) Non-Low to Non-Low

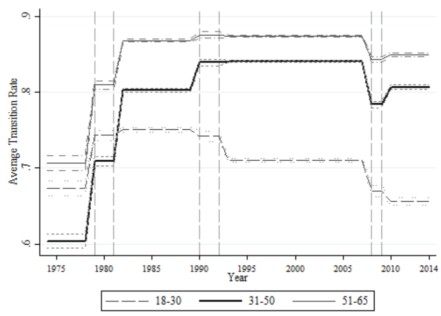

Period average transition rates from non-employment for female workers entering non-employment from low skilled job. Transition rates from low skilled non-employment to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.18: Non-employment outflow rates for female workers previously in medium skilled job

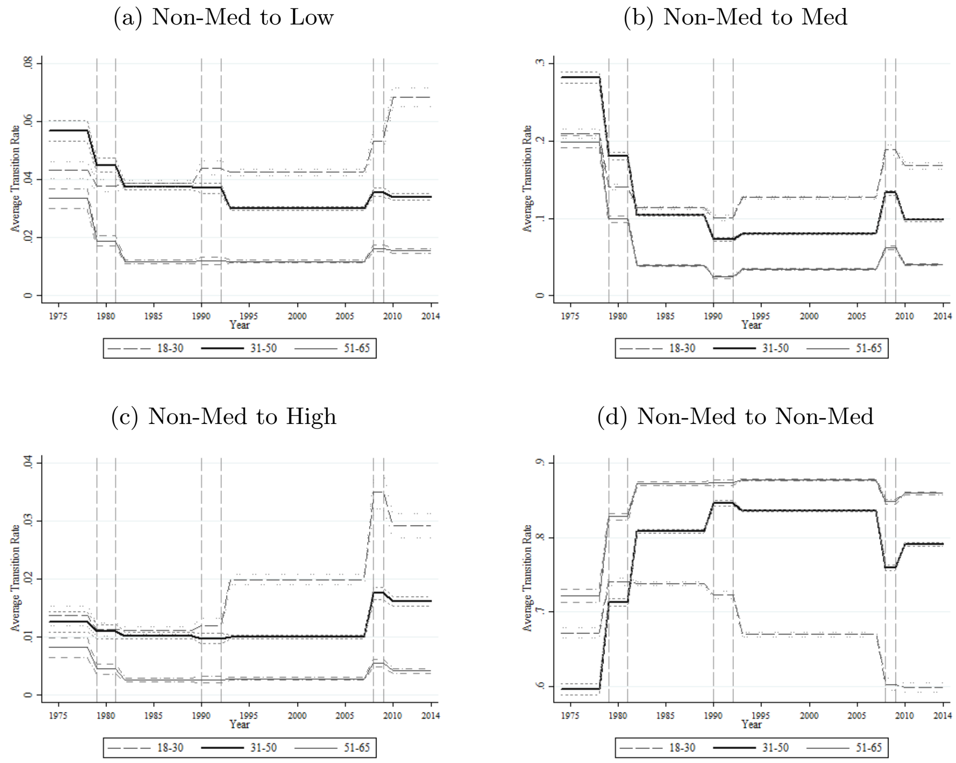

Period average transition rates from non-employment for female workers entering non-employment from medium skilled job. Transition rates from medium skilled non-employment to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.19: Non-employment outflow rates for female workers previously in high skilled job

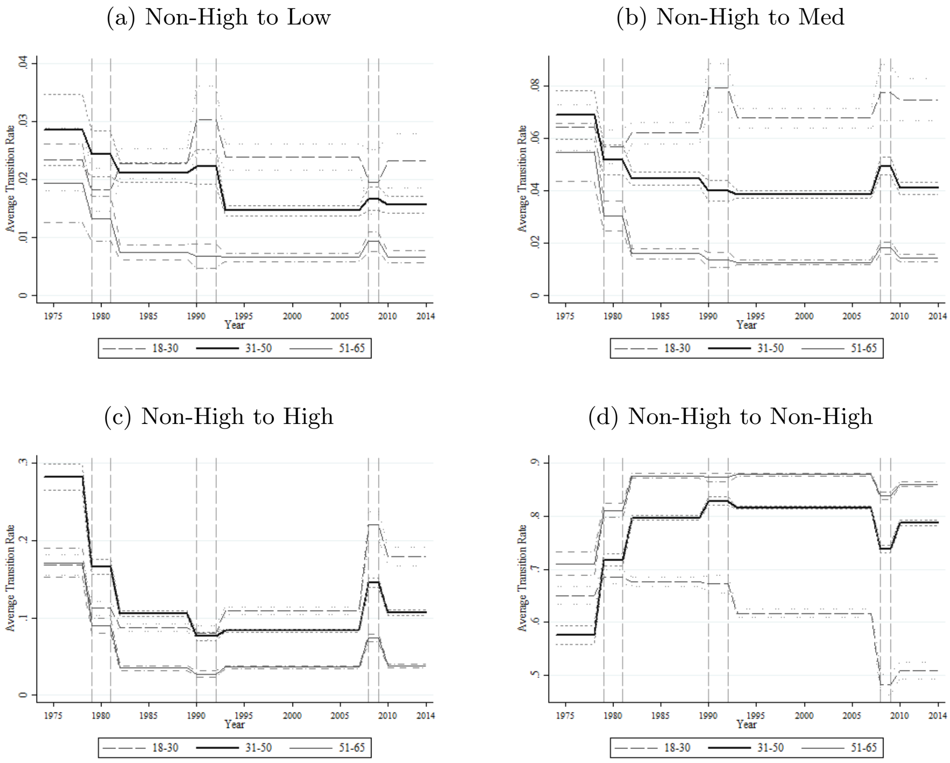

(a) Non-High to Low (b) Non-High to Med

(c) Non-High to High (d) Non-High to Non-High

Period average transition rates from non-employment for female workers entering non-employment from high skilled job. Transition rates from high skilled non-employment to low, medium, high skilled or non-employment are estimated based on regression equation 2.5, using a linear probability model with robust standard errors. Estimates correspond to period average transition rates. 95 percent confidence intervals are shown. Three age groups are considered (18-30, 31-50, 51-65). Years with discontinuities have been omitted: 1990, 1996, 2001, 2006, 2010. Recessionary periods are indicated by vertial lines: 1979-81, 1990-92, 2008-09. Based on annual observations of NESPD and ASHE.

Figure A.20: LFS population shares for male workers, 1975 to 2015

(a) Young worker

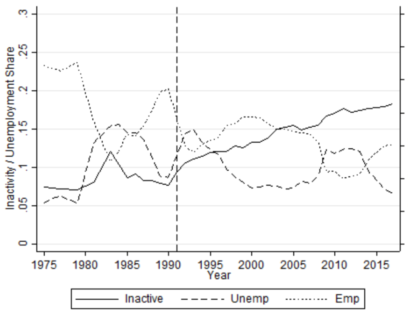

(b) Prime aged workers

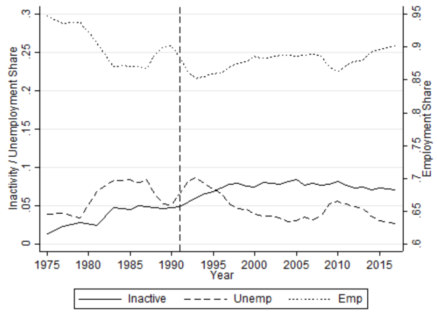

(c) Older workers

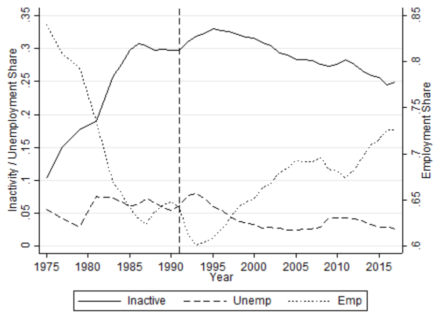

Annual population shares for inactive, unemployed, and employed male workers based on UK LFS from 1975 to 2015. Population corresponds to sum of inactive, unemployed and employed workers. 1975 to 1991 from annual LFS. 1992 to 2015 average annual population shares from quarterly LFS. Young workers are aged 18-30, prime aged workers 31-50, and older workers 51-65. Vertical line indicates change in 1991 from annual to quarterly LFS.

Figure A.21: LFS population shares for female workers, 1975 to 2015

(a) Young worker

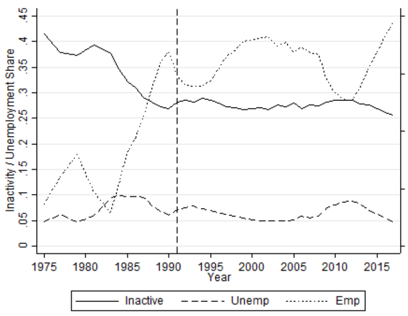

(b) Prime aged workers

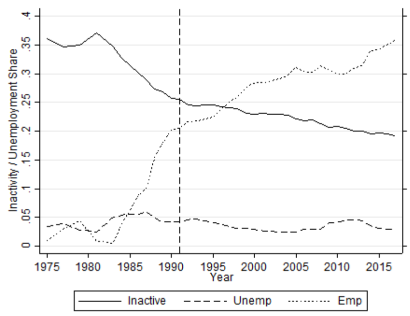

(c) Older workers

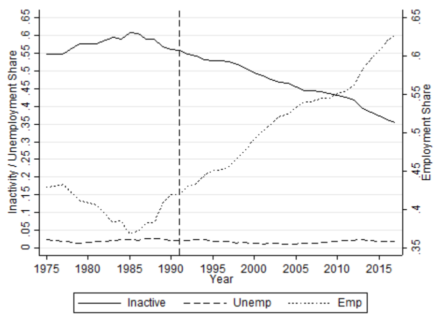

Annual population shares for inactive, unemployed, and employed female workers based on UK LFS from 1975 to 2015. Population corresponds to sum of inactive, unemployed and employed workers. 1975 to 1991 from annual LFS. 1992 to 2015 average annual population shares from quarterly LFS. Young workers are aged 18-30, prime aged workers 31-50, and older workers 51-65. Vertical line indicates change in 1991 from annual to quarterly LFS.

Appendix B

(For Chapter 3)

B.1 Derivation of equation 1

The derivation closely follows the general case discussed in Jenkins [2005]. To relate the overall interval hazard rate to destination-specific interval hazard rates, start from the continuous time hazard rate for duration t, defined as the limit of the probability to leave over the interval $t + \Delta t$, conditional on surviving until t:

$$\theta(t) = \lim_{\Delta t \to 0} \frac{\text{Prob}(t \leq T \leq t + \Delta t \mid T \geq t)}{\Delta t} = \frac{f(t)}{S(t)}$$

where $f(t) = \partial F/\partial t$. The discrete time interval hazard rate is related to the continuous time hazard rate as follows:

$$\lambda_j = 1 - \exp\biggl(- \int_{a_{j-1}}^{a_j} \theta(t) \mathrm{d}t \biggr)$$

Accordingly, the continuous time hazard rate for duration t to destination A is:

$$\theta^A(t) = \frac{f^A(t)}{S^A(t)}$$

where $A = L, M, H$. The discrete time destination-specific interval hazard rate is related to the continuous time destination-specific hazard rate as follows:

$$\lambda_j^A = 1 - \exp\biggl(- \int_{a_{j-1}}^{a_j} \theta^A(t) \mathrm{d}t \biggr)$$

Recall that $\lambda_j^A = 1 - S^A(a_j)/S^A(a_{j-1})$. It can be shown that the continuous time overall hazard rate is the sum of continuous time destination-specific hazard rates, i.e. $\theta(t) = \theta^L(t)\theta^M(t)\theta^H(t)$.[1] Using the fact that $S(a_j) = \exp\big(-\int_0^{a_j}\theta(t)\mathrm{d}t\big)$, one can express the overall interval hazard rate in terms of destination-specific interval hazard rates as follows:

$$\lambda_j = 1 - \frac{\exp\big(-\int_0^{a_j}\theta(t)\mathrm{d}t\big)}{\exp-\big(\int_0^{a_j-1}\theta(t)\mathrm{d}t\big)}$$

$$\lambda_j = 1 - \frac{\exp\big(-\int_0^{a_j}[\theta^L(t)+\theta^M(t)+\theta^H(t)]\mathrm{d}t\big)}{\exp\big(-\int_0^{a_j-1}[\theta^L(t)+\theta^M(t)+\theta^H(t)]\mathrm{d}t\big)}$$

$$\lambda_j = 1 - \exp\bigg(-\int_{a_{j-1}}^{a_j}[\theta^L(t)+\theta^M(t)+\theta^H(t)]\mathrm{d}t\bigg)$$

$$\lambda_j = 1 - \exp\bigg(-\int_{a_{j-1}}^{a_j}\theta^L(t)\mathrm{d}t\bigg)\exp\bigg(-\int_{a_{j-1}}^{a_j}\theta^M(t)\mathrm{d}t\bigg)\exp\bigg(-\int_{a_{j-1}}^{a_j}\theta^H(t)\mathrm{d}t\bigg)$$

$$1-\lambda_j = \Big(1-\lambda_j^L\Big)\Big(1-\lambda_j^M\Big)\Big(1-\lambda_j^H\Big)$$

[1] See Jenkins [2005], page 98.

B.2 Tables

Table B.1: Additional Decomposition of Survival Functions for Male Workers

Young Workers						
		Total	Contribution			
Period	Duration	Change	Low	Med	High	Residual
3	1	8.0	-0.3	6.0	0.8	1.5
	3	7.9	-0.4	6.4	0.7	1.1
	5	5.6	-0.7	4.8	0.6	0.9
	7	3.7	-0.8	3.4	0.4	0.7
	10	2.6	-0.9	2.5	0.3	0.7
5	1	12.9	-0.8	10.7	0.9	2.1
	3	9.4	-1.3	9.2	0.3	1.2
	5	5.2	-1.7	6.3	-0.2	0.8
	7	2.2	-2.0	4.2	-0.7	0.6
	10	0.4	-2.1	2.9	-0.9	0.5
7	1	10.8	-2.5	11.0	0.9	1.4
	3	8.0	-3.1	10.0	0.4	0.7
	5	4.2	-3.5	7.4	-0.1	0.4

Prime Aged Workers						
		Total	Contribution			
Period	Duration	Change	Low	Med	High	Residual
3	1	6.1	-0.1	4.7	0.2	1.3
	3	7.2	0.0	5.7	0.4	1.1
	5	5.5	-0.1	4.4	0.3	0.9
	7	4.2	-0.2	3.5	0.1	0.8
	10	3.3	-0.3	2.8	0.0	0.8
5	1	11.2	0.2	8.8	0.0	2.3
	3	5.9	-0.2	6.1	-1.0	1.1
	5	2.4	-0.5	3.6	-1.5	0.8
	7	0.0	-0.7	2.0	-1.9	0.7
	10	-1.8	-0.9	0.7	-2.3	0.6
7	1	10.0	-0.3	9.4	-0.8	1.7
	3	7.0	-0.7	8.5	-1.8	0.9
	5	3.7	-1.0	6.3	-2.2	0.7

Older Workers						
		Total	Contribution			
Period	Duration	Change	Low	Med	High	Residual
3	1	11.8	0.8	8.6	0.7	1.7
	3	14.4	1.1	10.5	1.1	1.6
	5	13.7	1.0	10.1	1.1	1.5
	7	13.1	0.9	9.7	1.1	1.5
	10	12.7	0.8	9.4	1.0	1.4
5	1	10.9	0.5	9.1	-0.2	1.4
	3	6.3	0.2	6.4	-1.2	0.9
	5	3.5	-0.2	4.6	-1.6	0.8
	7	1.6	-0.5	3.3	-1.9	0.7
	10	-0.2	-0.7	2.0	-2.2	0.7
7	1	5.9	0.2	7.0	-1.7	0.4
	3	2.5	-0.2	5.5	-2.8	0.1
	5	-0.9	-0.7	3.3	-3.5	0.0

Decomposition of changes in survival function relative to period 1975-78 at durations 1, 3, 5, 7, and 10. Decomposition based on equation 3.5. Survival functions based on Kaplan-Meier estimates for male workers entering non-employment in expansionary periods (1975-78, 1982-89, 2010-15). By age group: 18-30, 31-50, 51-65 years. Columns 4-7 give contribution of changes in hazard functions to low, medium, high skilled employment and decomposition residual. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Table B.2: Additional Decomposition of Survival Functions for Female Workers

Young Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
	1	-0.3	-0.4	0.1	0.1	-0.1
	3	-2.6	-1.7	-0.6	-0.1	-0.3
3	5	-5.1	-2.6	-1.9	-0.2	-0.4
	7	-6.9	-3.4	-2.7	-0.3	-0.5
	10	-6.1	-3.4	-2.2	-0.1	-0.4
	1	0.3	-0.8	2.0	-0.7	-0.2
	3	-7.1	-3.0	-1.2	-1.9	-1.0
5	5	-11.5	-4.1	-3.6	-2.5	-1.3
	7	-12.8	-4.6	-4.3	-2.7	-1.2
	10	-10.4	-3.9	-3.2	-2.3	-1.0
	1	-6.9	-4.9	1.5	-1.4	-2.2
7	3	-16.9	-9.3	-1.8	-2.7	-3.3
	5	-23.6	-11.9	-4.9	-3.2	-3.5

Prime Aged Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
	1	4.3	1.4	2.0	-0.3	1.2
	3	4.1	1.5	2.2	-0.5	1.0
3	5	1.9	0.8	1.1	-0.7	0.7
	7	0.4	0.1	0.4	-0.8	0.6
	10	-0.3	-0.3	0.2	-0.8	0.6
	1	6.5	2.8	2.8	-0.8	1.7
	3	0.4	1.5	0.2	-1.8	0.5
5	5	-3.1	0.3	-1.4	-2.3	0.2
	7	-5.6	-0.8	-2.4	-2.5	0.1
	10	-7.2	-1.7	-3.0	-2.6	0.0
	1	-1.5	1.0	0.4	-2.2	-0.8
7	3	-5.6	-0.4	-0.8	-3.2	1.3
	5	-8.6	-1.7	-2.1	-3.4	-1.4

Older Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
	1	9.1	2.9	4.3	0.3	1.6
	3	11.3	3.8	5.6	0.4	1.6
3	5	11.0	3.7	5.4	0.4	1.5
	7	10.5	3.4	5.4	0.3	1.5
	10	10.3	3.1	5.4	0.3	1.4
	1	6.6	2.3	3.7	-0.6	1.2
	3	1.2	0.7	1.5	-1.6	0.6
5	5	-0.9	0.0	0.4	-1.8	0.5
	7	-2.3	-0.7	-0.2	-1.9	0.5
	10	-3.3	-1.4	-0.4	-2.1	0.5
	1	-2.0	-0.5	0.6	-1.4	-0.7
7	3	-7.0	-2.4	-1.3	-2.3	-1.0
	5	-9.8	-3.7	-2.5	-2.6	-1.0

Decomposition of changes in survival function relative to period 1975-78 at durations 1, 3, 5, 7, and 10. Decomposition based on 3.5. Survival functions based on Kaplan-Meier estimates for female workers entering non-employment in expansionary periods (1975-78, 1982-89, 2010-15). By age group: 18-30, 31-50, 51-65 years. Columns 4-7 give contribution of changes in hazard functions to low, medium, high skilled employment and decomposition residual. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Table B.3: Decomposition of Survival Functions for Medium Skilled Male Workers

Young Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
3	1	8.0	-0.2	7.5	0.0	0.7
	5	6.3	-0.4	6.2	0.1	0.5
	10	3.2	-0.7	3.7	-0.1	0.3
5	1	12.4	-0.2	11.8	-0.2	1.0
	5	5.2	-0.8	6.5	-0.8	0.3
	10	0.4	-1.2	2.8	-1.4	0.1
7	1	10.2	-0.9	10.9	-0.3	0.5
	5	4.3	-1.7	6.9	-0.8	0.0

Prime Aged Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
3	1	6.9	-0.2	7.3	-0.5	0.3
	5	6.2	-0.2	6.6	-0.4	0.3
	10	4.0	-0.3	4.6	-0.5	0.2
5	1	11.6	0.1	11.3	-0.5	0.9
	5	2.7	-0.4	4.3	-1.3	0.1
	10	-1.6	-0.7	0.8	-1.8	0.1
7	1	9.6	-0.1	10.8	-1.3	0.2
	5	3.1	-0.7	6.5	-2.4	-0.2

Older Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
3	1	12.0	0.1	11.2	-0.1	0.7
	5	13.8	0.2	12.7	0.3	0.6
	10	12.9	0.1	11.9	0.3	0.6
5	1	10.8	0.0	10.6	-0.3	0.5
	5	3.0	-0.6	4.4	-1.0	0.1
	10	-1.0	-1.1	1.4	-1.4	0.1
7	1	4.7	-0.1	6.5	-1.4	-0.3
	5	-3.1	-0.7	1.1	-2.9	-0.6

Decomposition of changes in survival function relative to period 1975-78 at durations 1, 5, and 10. Decomposition based on 3.5. Survival functions based on Kaplan-Meier estimates for male workers entering non-employment from medium skilled jobs in expansionary periods (1975-78, 1982-89, 2010-15). By age group: 18-30, 31-50, 51-65 years. Columns 4-7 give contribution of changes in hazard functions to low, medium, high skilled employment and decomposition residual. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Table B.4: Decomposition of Survival Functions for Medium Skilled Female Workers

Young Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
3	1	-0.9	-1.0	1.1	-0.5	-0.6
	5	-5.4	-2.4	-1.5	-0.7	-0.7
	10	-5.4	-2.4	-1.5	-0.7	-0.7
5	1	-1.7	-1.0	1.5	-1.3	-0.9
	5	-13.3	-2.5	-6.0	-3.1	-1.7
	10	-11.4	-2.0	-5.2	-2.9	-1.3
7	1	-9.7	-2.7	-2.2	-1.9	-2.9
	5	-25.3	-7.8	-9.7	-3.8	-3.9

Prime Aged Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
3	1	3.9	-0.7	5.3	-0.6	0.0
	5	1.8	-1.2	3.7	-0.7	0.0
	10	-0.4	-1.9	2.4	-0.9	-0.1
5	1	5.4	0.0	6.0	-0.9	0.3
	5	-4.7	-1.1	-1.2	-1.8	-0.5
	10	-8.7	-2.5	-3.4	-2.2	-0.6
7	1	-3.6	-0.5	0.7	-1.9	-1.8
	5	-9.9	-1.9	-3.2	-2.9	-1.9

Older Workers

Period	Duration	Total Change	Contribution Low	Med	High	Residual
3	1	9.3	0.4	8.1	0.0	0.7
	5	11.0	0.5	9.8	0.0	0.7
	10	9.9	0.0	9.3	0.0	0.7
5	1	6.6	-0.3	6.6	-0.1	0.4
	5	-0.9	-1.3	1.0	-0.7	0.1
	10	-3.7	-2.4	-0.4	-0.9	0.0
7	1	-3.0	-1.5	0.4	-0.8	-1.1
	5	-12.0	-4.2	-4.7	-1.7	-1.4

Decomposition of changes in survival function relative to period 1975-78 at durations 1, 5, and 10. Decomposition based on 3.5. Survival functions based on Kaplan-Meier estimates for female workers entering non-employment from medium skilled jobs in expansionary periods (1975-78, 1982-89, 2010-15). By age group: 18-30, 31-50, 51-65 years. Columns 4-7 give contribution of changes in hazard functions to low, medium, high skilled employment and decomposition residual. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

B.3 Figures

Figure B.1: Observed Versus Predicted Survival Functions for Male Workers

(a) Young workers

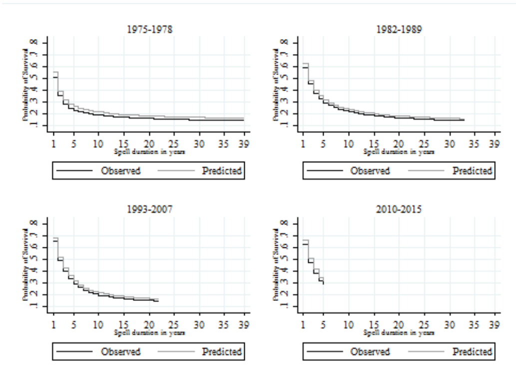

(b) Prime aged workers

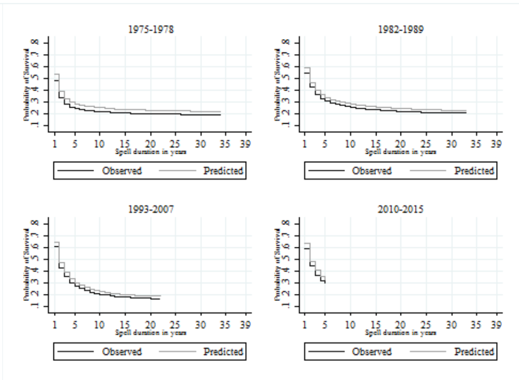

(c) Older workers

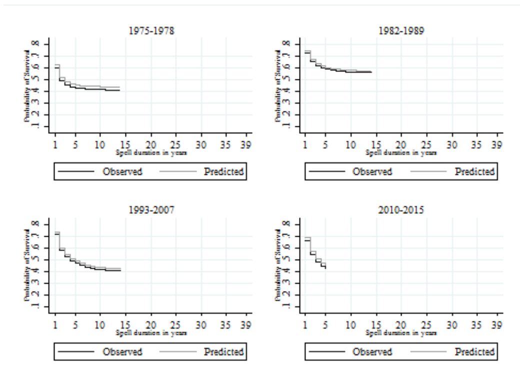

Comparing predicted survival function assuming independence of competing risks, with observed survival functions. Kaplan-Meier estimates for survival functions to employment for male workers entering non-employment in expansionary periods: 1975-78 (period 1), 1982-89 (period 3), 1993-07 (period 5), 2010-15 (period 7). By age group: 18-30, 31-50, 51-65 years. For each group and period, observed survival function is computed according to equation 3.2, predicted according to equation 3.4. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure B.2: Observed Versus Predicted Survival Functions for Female Workers

(a) Young workers

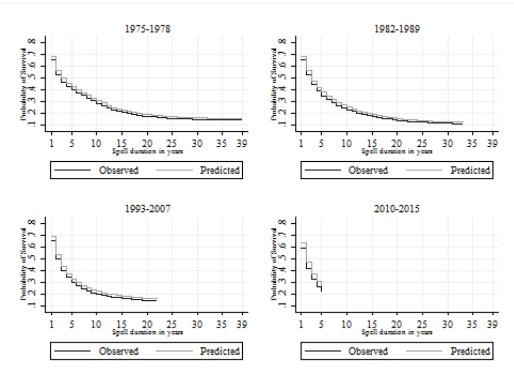

(b) Prime aged workers

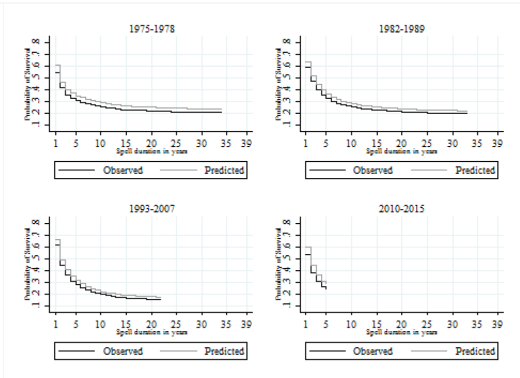

(c) Older workers

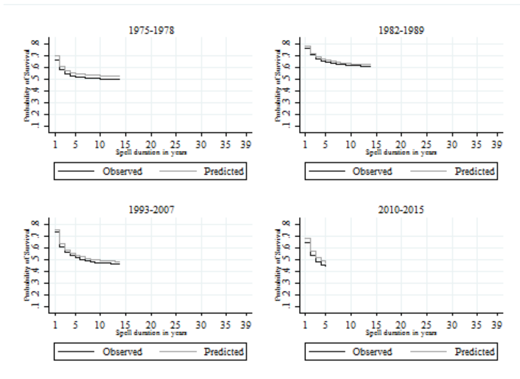

Comparing predicted survival function assuming independence of competing risks, with observed survival functions. Kaplan-Meier estimates for survival functions to employment for female workers entering non-employment in expansionary periods: 1975-78 (period 1), 1982-89 (period 3), 1993-07 (period 5), 2010-15 (period 7). By age group: 18-30, 31-50, 51-65 years. For each group and period, observed survival function is computed according to equation 3.2, predicted according to equation 3.4. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure B.3: Hazard Functions to Low Skilled Jobs for Male Workers

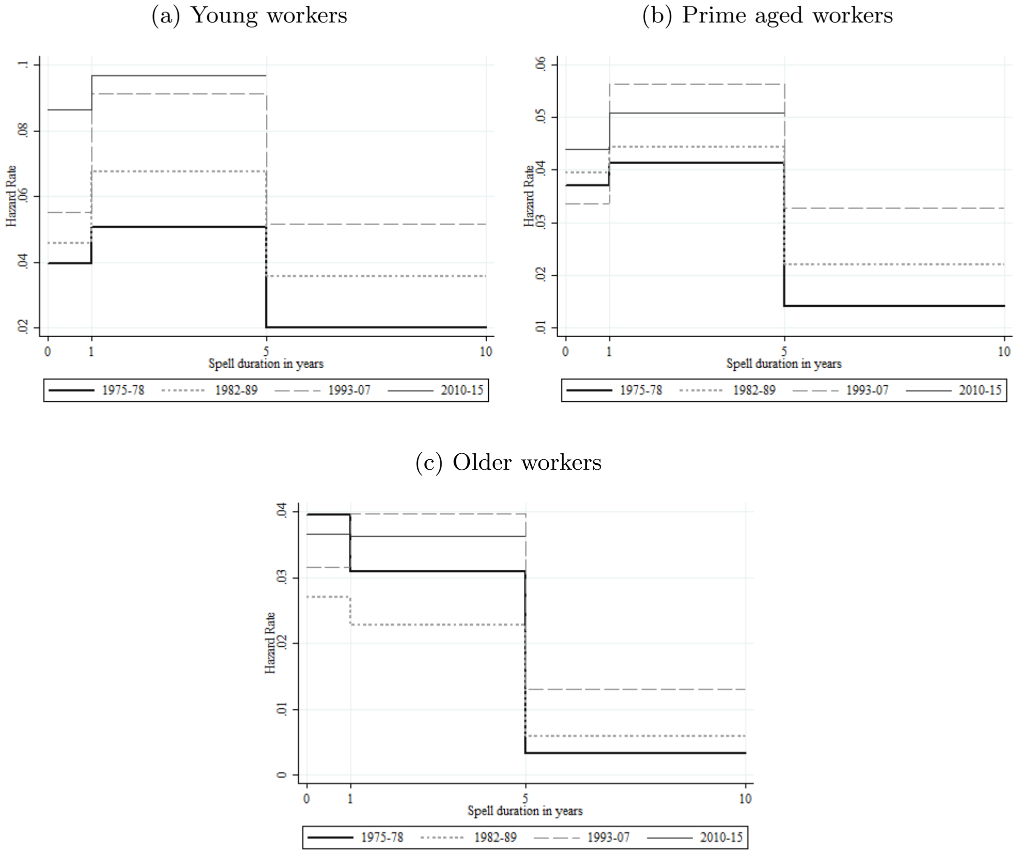

Kaplan-Meier estimates for interval hazard function to low skilled employment for male workers entering non-employment in expansionary periods: 1976-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.1. Interval 0-1 gives hazard rate at duration 1, 1-5 at durations 2-4, 5-10 at durations 5 or longer. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure B.4: Hazard Functions to Medium Skilled Jobs for Male Workers

(a) Young workers

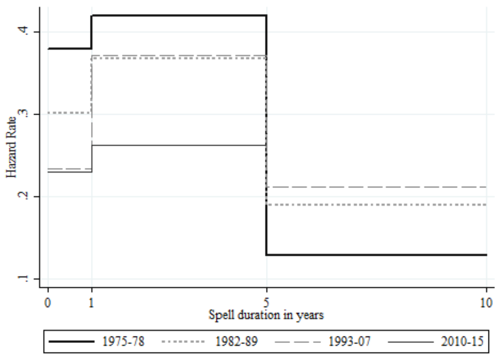

(b) Prime aged workers

(c) Older workers

Kaplan-Meier estimates for interval hazard function to medium skilled employment for male workers entering non-employment in expansionary periods: 1976-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.1. Interval 0-1 gives hazard rate at duration 1, 1-5 at durations 2-4, 5-10 at durations 5 or longer. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure B.5: Hazard Functions to High Skilled Jobs for Male Workers

(a) Young workers

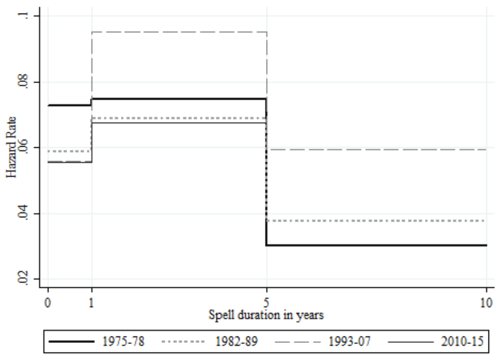

(b) Prime aged workers

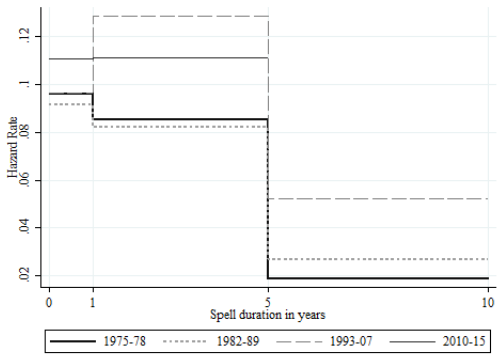

(c) Older workers

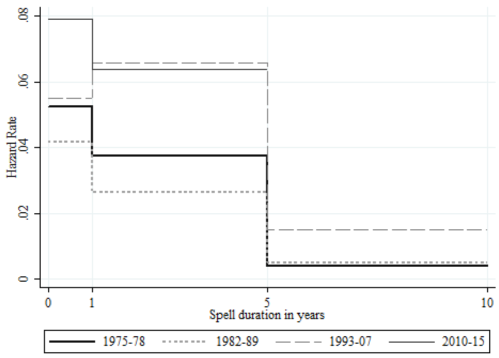

Kaplan-Meier estimates for interval hazard function to high skilled employment for male workers entering non-employment in expansionary periods: 1976-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.1. Interval 0-1 gives hazard rate at duration 1, 1-5 at durations 2-4, 5-10 at durations 5 or longer. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure B.6: Hazard Functions to Low Skilled Jobs for Female Workers

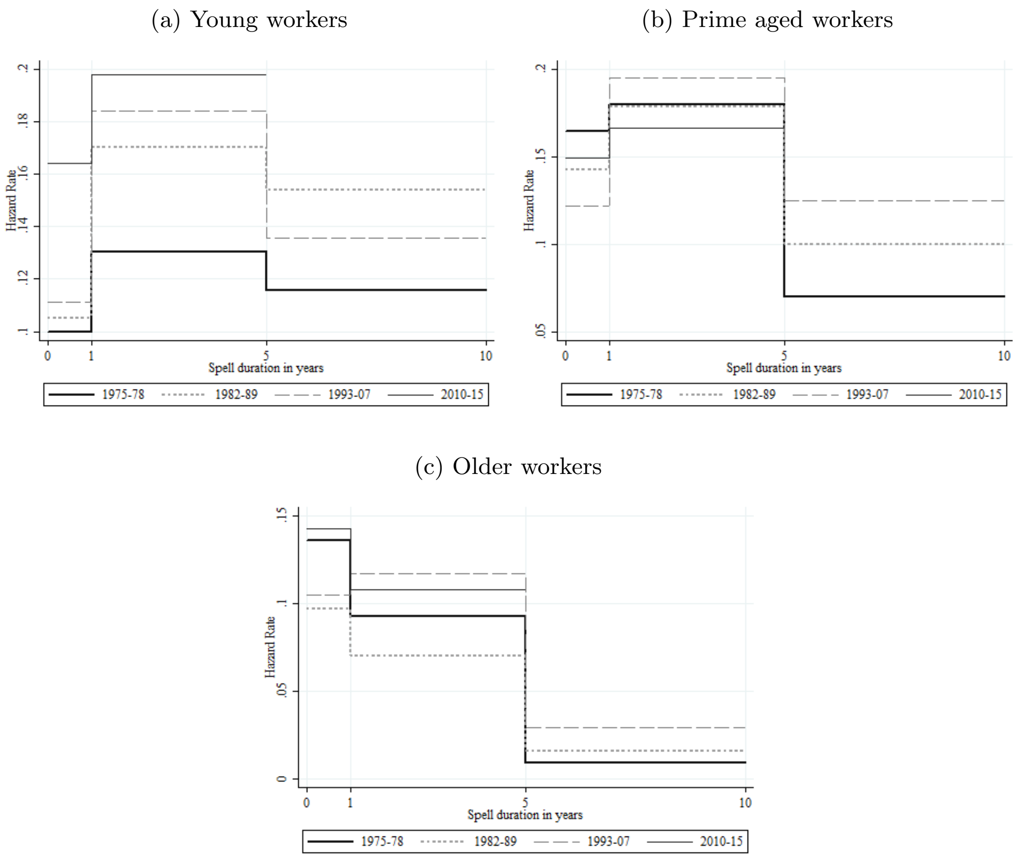

Kaplan-Meier estimates for interval hazard function to low skilled employment for female workers entering non-employment in expansionary periods: 1976-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.1. Interval 0-1 gives hazard rate at duration 1, 1-5 at durations 2-4, 5-10 at durations 5 or longer. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure B.7: Hazard Functions to Medium Skilled Jobs for Female Workers

(a) Young workers

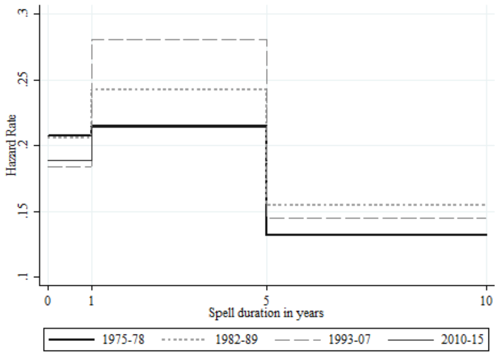

(b) Prime aged workers

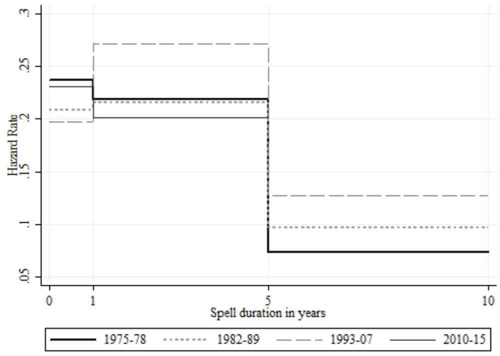

(c) Older workers

Kaplan-Meier estimates for interval hazard function to medium skilled employment for female workers entering non-employment in expansionary periods: 1976-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.1. Interval 0-1 gives hazard rate at duration 1, 1-5 at durations 2-4, 5-10 at durations 5 or longer. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure B.8: Hazard Functions to High Skilled Jobs for Female Workers

(a) Young workers

(b) Prime aged workers

(c) Older workers

Kaplan-Meier estimates for interval hazard function to high skilled employment for female workers entering non-employment in expansionary periods: 1976-78, 1982-89, 2010-15. By age group: 18-30, 31-50, 51-65 years. Estimated for each group and period according to equation 3.1. Interval 0-1 gives hazard rate at duration 1, 1-5 at durations 2-4, 5-10 at durations 5 or longer. Periods divide expansionary from recessionary periods in 1979-81, 1990-92, 2008-09. Based on annual observations from NESPD and ASHE.

Figure B.9: LFS population shares for male workers, 1975 to 2015

(a) Young worker

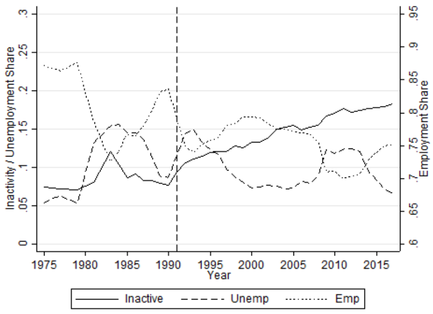

(b) Prime aged workers

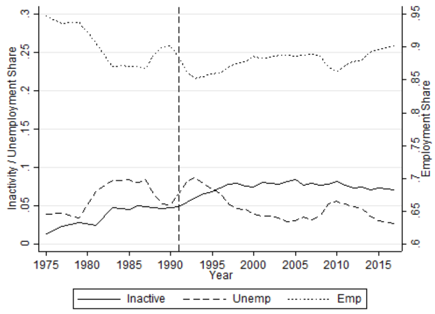

(c) Older workers

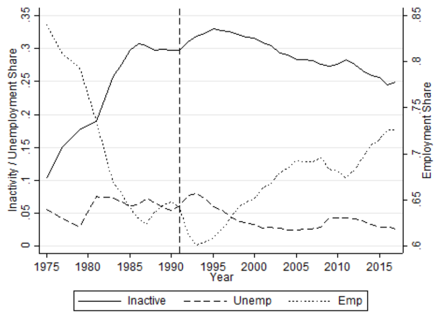

Annual population shares for inactive, unemployed, and employed male workers based on UK LFS from 1975 to 2015. Population corresponds to sum of inactive, unemployed and employed workers. 1975 to 1991 from annual LFS. 1992 to 2015 average annual population shares from quarterly LFS. Young workers are aged 18-30, prime aged workers 31-50, and older workers 51-65. Vertical line indicates change in 1991 from annual to quarterly LFS.

Figure B.10: LFS population shares for female workers, 1975 to 2015

(a) Young worker

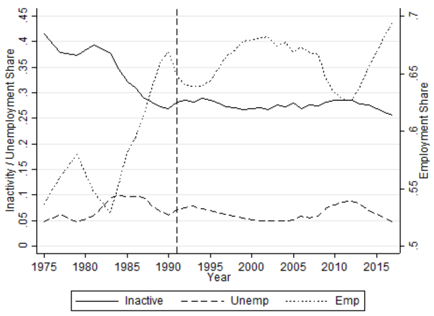

(b) Prime aged workers

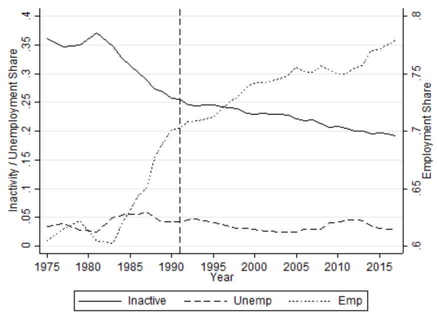

(c) Older workers

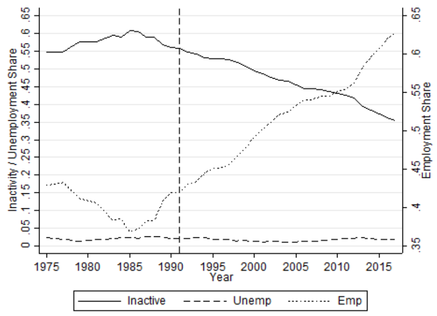

Annual population shares for inactive, unemployed, and employed female workers based on UK LFS from 1975 to 2015. Population corresponds to sum of inactive, unemployed and employed workers. 1975 to 1991 from annual LFS. 1992 to 2015 average annual population shares from quarterly LFS. Young workers are aged 18-30, prime aged workers 31-50, and older workers 51-65. Vertical line indicates change in 1991 from annual to quarterly LFS.

Appendix C

(For Chapter 4)

C.1 Data

All data refer to annual observations at the country level, where applicable for men and women. The particular country coverage varies by variable. The countries covered in the regression analyses are Australia, Denmark, Finland, France, Germany (West-Germany before 1991), Ireland, Italy, Japan, Korea, Norway, Sweden, Switzerland, United Kingdom, and USA.

Relative skill supply My measure for relative skill supply is based on educational attainment data taken from taken from Barro and Lee [2013]. Barro and Lee [2013] use census and survey data and fill in missing data in 5 year intervals by extrapolating existing observations backward and forward for various age groups based on assumptions about mortality rates. I construct a variable for relative skill supply as the ratio of the number of people as percentage of the population with at least some tertiary education over the sum of the number of people as percentage of the population with at least some primary and secondary education. As the denominator of each percentage term cancels out, the resulting variable gives the number of people with some tertiary education, which I take to represent skilled workers, over the sum of the number of people with either some primary and some secondary education, which I take to represent unskilled workers. As the data in 5 year intervals exhibits constant trends, I extrapolate the data in 5 year intervals to attain yearly observations.

Earnings decile ratio The earnings decile ratio is the ratio of 9th-to-1st upper-earnings deciles of gross earnings of full-time dependent employees. Data are taken from OECD [2015] (DOI: 10.1787/lfs-ear-data-en).

Temporary and regular employment protection Strictness of employment protection is a synthetic measure ranging from 1 to 6 with higher values indicating stricter protection. There exists a separate measure for temporary and regular contracts each, in both cases based on several items.

For temporary contracts, the measure is constructed based on regulations for temporary work agencies and for temporary workers in relation to regular workers. For regular contracts, the measure combines items on individual, comprising measures on the notification time prior to dismissals, severance payments, as well as legal issues surrounding the possibility to dismiss workers and the possibility and costs related to to legally disputing dismissals. There exist two other versions of OECD data incorporating more items and additionally taking into account regulation on collective dismissals. However, these datasets only cover the years 1998 to 2013, so I restrict my attention to the more narrow definition covering data from 1985 onwards. Data are taken from OECD [2015] (DOI: 10.1787/lfs-epl-data-en).

Unemployment benefit net replacement rate and unemployment benefit duration The unemployment benefit net replacement rate is based on two measures. Measures indicate the generosity of unemployment benefits in terms of the benefits received during unemployment as percentage of previous earnings. Both measures are based on an unweighted average of net replacement rates received in situations varying in terms of family status, duration of unemployment and previous earnings relative to some average benchmark worker. The measures are distinct in that they consider different benchmark workers.

Data from 1982 until 2003 are based on earnings relative to an average 40 year old full-time production worker in the manufacturing sector with a long previous employment history (APW). Data from 2001 until 2010 are based on earnings relative to an average 40 year old full-time worker in the private sector with a long previous employment history (AW). As the measure based on APW was discontinued, I use a combination of both measures with data based on APW until 2000 and based on AW from 2001 until 2010. To account for the break, in any regression using this variable I include a dummy variable which indicates whether data based on AW or APW were used. Unemployment benefit duration is calculated as the initial replacement rate over the average replacement rate, following the method explained above. Data for both measures are taken from the ifo Institute [2013] database, which is based on OECD data.

Wage setting coordination The variable on coordination of wage setting is a categorical variable ranging from 1 to 5, with higher values indicating more coordination. The measure reflects both formal and informal arrangements as well as de facto coordination, and may thus reflect high levels of coordination even if bargaining takes place at more decentralized levels. Data are taken from Visser [2015].

Labour force participation rate The labour force participation rate for the age group 15 to 64 is expressed in terms of the total labour force as percentage of the population. Data are taken from OECD [2015] (DOI : 10.1787/8a801325-en).

Trade openness The measure on trade openness is calculated as the sum of imports plus exports divided by GDP. Data for the three series are taken from OECD [2015] (DOI: 10.1787/data-00285-en). Imports and Exports in goods, current prices, billions USD, GDP current prices current exchange rates, millions USD.

Additional Control Variables The employment share of the service sector is computed as the fraction of total employment in services of total employment. Total employment and total employment in services in thousands are taken from the OECD [2015] (doi: 10.1787/a258bb52-en). Relative educational expenditures are computed as the ratio of total expenditures on teritary over total expenditures on primary education. Total expenditures are measured as fraction of GDP. Total expenditures comprise public and private expenditures. Data are taken for educational expenditures are taken from Brady et al. [2014]. Data are extrapolated.

C.2 Tables

Table C.1: Alternative Wage Inequality Measures

	P90P50	P50P10	P90P50	P50P10	P90P50	P50P10
RSS	0.0842** (4.02)	0.0306 (1.34)	-0.00357 (-0.06)	-0.0687 (-1.46)	-0.128 (-0.95)	-0.000390 (-0.00)
RSS Can	-0.0521* (-2.86)	-0.0535 (-1.99)			-0.0140 (-0.36)	-0.0818 (-2.05)
RSS Den	-0.0689** (-3.66)	0.223*** (17.68)			-0.0503 (-1.30)	0.221*** (5.93)
RSS Fin	-0.0649*** (-4.25)	-0.0475* (-2.49)			-0.0197 (-0.80)	-0.00434 (-0.20)
RSS Fra	-0.124*** (-6.80)	-0.142*** (-7.47)			-0.0837** (-3.54)	-0.0863* (-2.96)
RSS Ger	-0.140 (-1.96)	0.159 (1.74)			-0.0964 (-1.63)	0.204 (2.06)
RSS Ire	-0.0959** (-3.03)	-0.111* (-2.29)			-0.148* (-2.61)	-0.192* (-2.29)
RSS Ita	-0.0583 (-0.90)	-0.0849 (-1.81)			-0.231*** (-5.04)	-0.0325 (-0.46)
RSS Jap	-0.134*** (-13.74)	-0.122*** (-10.91)			-0.130*** (-7.41)	-0.152*** (-5.84)
RSS Kor	0.131*** (5.91)	-0.0475 (-1.20)			0.153*** (4.61)	0.00538 (0.09)
RSS Nrw	0.108 (1.79)	0.339*** (6.73)			0.129* (2.26)	0.409*** (5.72)
RSS Swe	0.0743 (1.53)	0.000876 (0.02)			0.124* (2.82)	0.0267 (0.53)
RSS Swz	0.0685** (3.49)	-0.0406** (-3.15)			0.0763*** (4.47)	-0.0528** (-4.13)
RSS UK	0.0296 (1.85)	-0.0635** (-3.53)			0.0786 (1.89)	-0.0683 (-1.74)
RSS USA	0.189** (4.10)	-0.00243 (-0.06)			0.0962 (1.24)	-0.123 (-1.72)
Inst 1	-0.0181 (-0.64)	-0.00769 (-0.34)	0.0623 (1.11)	0.00905 (0.14)	0.0694 (1.31)	-0.0588 (-0.71)
Inst 2	0.0111 (0.28)	0.0183 (0.84)	0.123** (3.88)	0.0511 (1.66)	0.140 (1.75)	0.0803 (1.15)
Inst 1*2	-0.000801 (-0.06)	-0.00306 (-0.46)	-0.0497** (-4.11)	-0.0130 (-0.81)	-0.0512* (-2.26)	-0.0217 (-1.11)
RSS Inst 1			0.0239 (1.43)	0.00937 (0.44)	0.0679* (2.23)	-0.00975 (-0.26)
RSS Inst 2			0.0487 (1.60)	0.0308 (1.21)	0.0935 (1.61)	0.0620 (1.08)
RSS Inst 1*2			-0.0213* (-2.80)	-0.00758 (-1.18)	-0.0339* (-2.35)	-0.0178 (-1.34)
Emp/Pop	0.00240 (1.69)	-0.000701 (-0.28)	0.00129 (0.71)	-0.00115 (-0.51)	0.00352* (2.37)	0.000329 (0.15)
Trade	-0.0265 (-0.53)	-0.0545 (-0.64)	0.0928* (2.15)	0.0375 (0.63)	-0.0349 (-0.70)	-0.0734 (-0.85)
Fixed Effects	Yes	Yes	No	No	Yes	Yes
Period	1986-2010	1986-2010	1986-2010	1986-2010	1986-2010	1986-2010
N	298	298	298	298	298	298

Note: t statistics in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Results for regression specification 4.4 in columns 1 and 2, specification 4.5 in columns 3 and 4, and specficiation 4.6 in columns 5 and 6. P90P50 is the log of the 9th to 5th earnings decile, P50P10 is the log of the 5th to th 1st decile. 'RSS' is the log of relative skill supply and measures the common slope. Deviations from common slope are given by country-suffix RSS. 'Inst 1' and 'Inst 2' refer to factor variables 1 and 2, extracted from institutional variables. Details for factor variables given in 4.3.1. Using clustered standard errors at country-level. Regression includes constant. Slope deviation for Australia not shown.

Table C.2: Additional Control Variables

	Control 1	Control 2	Control 1	Control 2	Control 1	Control 2
RSS	0.0548 (1.62)	0.0333 (1.41)	-0.0960 (-1.25)	-0.0896 (-1.29)	-0.200 (-1.96)	-0.283* (-3.19)
RSS Den	0.169*** (7.70)	0.207*** (10.60)			0.236*** (7.06)	0.218*** (8.43)
RSS Fin	-0.0584 (-2.31)	-0.0339 (-1.40)			0.0348 (1.12)	0.0477 (2.15)
RSS Fra	-0.283*** (-11.09)	-0.216*** (-14.92)			-0.190*** (-6.72)	-0.117** (-4.21)
RSS Ger	-0.0869 (-1.09)	0.0216 (0.22)			0.0104 (0.11)	0.117 (0.97)
RSS Ita	-0.0999 (-1.20)	-0.126 (-1.46)			-0.312** (-4.41)	-0.249** (-3.99)
RSS UK	-0.00569 (-0.31)	-0.0203 (-1.07)			0.0411 (1.10)	0.0614 (2.13)
RSS USA	0.233** (4.02)	0.0629 (1.07)			0.0853 (1.43)	-0.124 (-1.62)
Inst 1	0.0150 (0.42)	0.0423 (1.50)	-0.0304 (-0.41)	-0.0417 (-0.55)	0.144 (1.87)	0.113 (1.32)
Inst 2	0.0353 (0.74)	0.0571 (1.53)	0.204* (3.45)	0.171* (3.28)	0.222** (3.60)	0.264** (5.39)
Inst 1*2	-0.00530 (-0.29)	-0.0130 (-0.96)	-0.0583* (-2.92)	-0.0423* (-2.43)	-0.0879** (-4.19)	-0.0859** (-5.21)
RSS Inst 1			0.0106 (0.46)	-0.000461 (-0.02)	0.0873* (2.79)	0.0775 (2.34)
RSS Inst 2			0.153** (4.07)	0.134** (3.83)	0.119* (2.47)	0.166** (4.19)
RSS Inst 1*2			-0.0434** (-4.58)	-0.0345** (-4.50)	-0.0483** (-4.28)	-0.0532*** (-6.18)
Emp/Pop	0.00668** (3.60)	0.00547* (2.58)	0.00757** (4.68)	0.00688** (4.71)	0.00858*** (6.49)	0.00727*** (7.88)
Trade	0.119 (1.01)	0.106 (0.91)	0.124 (1.41)	0.133 (1.49)	0.110 (1.03)	0.0893 (0.80)
Edu Qual		-0.0548 (-2.06)		-0.0332** (-3.63)		-0.0536** (-3.87)
Fixed Effects	Yes	Yes	No	No	Yes	Yes
Period	1986-2010	1986-2010	1986-2010	1986-2010	1986-2010	1986-2010
N	175	175	175	175	175	175

Note: t statistics in parentheses.* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Results for regression specification 4.4 in columns 1 and 2, specification 4.5 in columns 3 and 4, and specficiation 4.6 in columns 5 and 6. Dependent variable is the log of the 9th to 1st earnings decile. 'Control 1' refers to the model in the main analysis, controlling for employment share and trade openness. 'Control 2' additionally controls for expenditures on teritary relative to primary education, 'Edu Qual'. 'RSS' is the log of relative skill supply and measures the common slope. Deviations from common slope are given by country-suffix RSS. 'Inst 1 'and 'Inst 2' refer to factor variables 1 and 2, extracted from institutional variables. Details for factor variables given in 4.3.1. Using clustered standard errors at country-level. Regression includes constant. Slope deviation for Australia not shown.

C.3 Figures

Figure C.1: Institutional Measures Across Countries

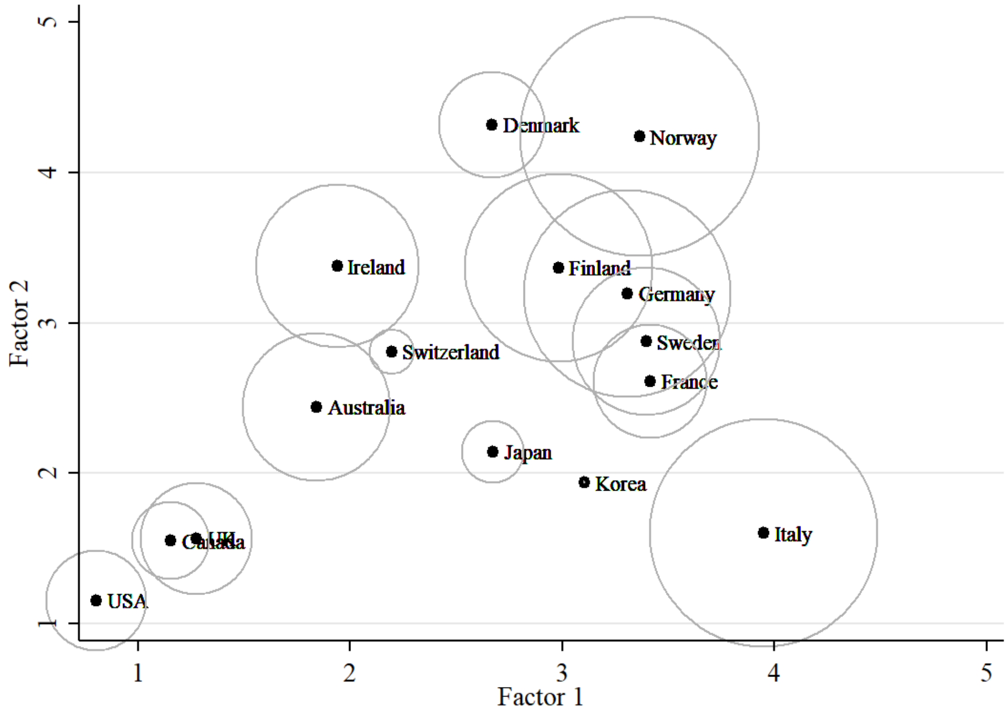

The figures shows the country-specific mean for factor variables 1 and 2. The size of the circle around the mean gives institutional variation, measured as the sum of variance for both factors. Factor variables range from 0 to 5. Higher values indicate more restrictive institutions. Details for factor variables are given in 4.3.1.

Bibliography

Daron Acemoglu. Cross-country inequality trends. *The Economic Journal*, 113(1): 121–149, 2003.

Daron Acemoglu. *Introduction to Modern Economic Growth*. Princeton Press, 2009.

Daron Acemoglu and David Autor. Skills, tasks and technologies: Implications for employment and earnings. In David Card Orley Ashenfelter, editor, *Handbook of Labor Economics*, volume 4b, pages 1043—-1171. Elsevier, 2011.

Daron Acemoglu and Fabrizio Zilibotti. Productivity differences. *The Quarterly Journal of Economics*, 116(2):563–606, 2001.

Philippe Aghion, Ufuk Akcigit, and Peter W. Howitt. What do we learn from schumpeterian growth theory? In Philipe Aghion and Steven N. Durlauf, editors, *Handbook of Economic Growth*, volume 2B, pages 515–563. Elsevier, 2014.

Paul D. Allison. *Event History and Survival Analysis*. Quantitative Applications in the Social Sciences. SAGE, 2 edition, 2014.

David Autor. Polanyi's paradox and the shape of employment growth. *NBER Working Paper Series*, (20485):1–46, 2014.

David Autor, Lawrence F. Katz, and Melissa S. Kearney. Measuring and interpreting trends in economic inequality - the polarization of the us labor market. *AEA Papers and Proceedings*, 96(2), 189-194.

David Autor, Lawrence F. Katz, and Alan B. Krueger. Computing inequality: Have computers chaned the labor market? *The Quarterly Journal of Economics*, 113 (4):1169–1213, 1998.

David Autor, Frank Levy, and Richard J. Murnane. The skill content of recent technological change: An empirical exploration. *The Quarterly Journal of Economics*, 118:1279–1333, 2003.

David Autor, David Dorn, and Gordon H. Hanson. Untangling trade and technology: Evidence from local labour markets. *The Economic Journal*, 125:621–646, 2015.

Dean Baker, Andrew Glyn, David R. Howell, and John Schmitt. Labor market institutions and unemployment: Assessment of the cross-country evidence. In David R. Howell, editor, *Fighting Unemployment: The Limits of Free Market Orthodoxy*, pages 72–117. Oxford University Press, 2005.

Robert Barro and Jong-Wha Lee. A new data set of educational attainment in the world, 1950-2010. *Journal of Development Economics*, 104:184–198, 2013.

Eric Bartelsman. Ict, reallocation and productivity. *European Commission Economic Papers*, (486):1–44, 2013.

Eli Berman and Stephen Machin. Skill-biased technology transferred around the world. *Oxford Review of Economic Policy*, 16(3), 2000.

Eli Berman, John Bound, and Zvi Griliches. 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, 1994.

Eli Berman, John Bound, and Stephen Machin. Implications of skill-biased technological change: International evidence. *The Quarterly Journal of Economics*, 113 (4):1245–1279, 1998.

Olivier Blanchard. European unemployment: The evolution of facts and ideas. *Economic Policy*, 1:5–59, 2006.

Olivier Blanchard and Justin Wolfers. The role of shocks and institutions in the rise of european unemployment: The aggregate evidence. *NBER Working Paper Series*, (7282):1–53, 1999.

John Bound and George Johnson. Changes in the structure of wages in the 1980s: An evalution of alternative explanations. *The American Economic Review*, 82(3): 371–392, 1992.

David Brady, Evelyne Huber, , and John D. Stephens. Comparative welfare states dataset, 2014.

Erik Brynjolfsson and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies.* W. W. Norton Company, 2014.

Pierre Cahuc, Stephane Carcillo, and Andre Zylberberg. *Labor Economics*. The MIT Press, 2 edition, 2014.

David Card and John E. DiNardo. Skill biased technological change and rising wage inequality: Some problems and puzzles. *NBER Working Paper Series*, (8769): 1–49, 2002.

David Card, Francis Kramarz, and Thomas Lemieux. Changes in the relative structure of wages and employment: A comparison of the united states, canada and france. *The Canadian Journal of Economics*, 32(4):842–877, 1999.

Francesco Caselli. Accounting for cross-country income differences. In Philippe Aghion and Steven N. Durlauf, editors, *Handbook of Economic Growth*, volume 1A, pages 680–741. Elsevier, 2005.

Francesco Caselli and Wilbur John Coleman II. The world technology frontier. *The American Economic Review*, 96(3):499–522, 2006.

Center of Economic Performance. Incapacity benefit reform: Tackling the rise in labour market inactivity. *CEP Policy Analysis*, pages 1–8, 2006.

R. H. Coase. The nature of the firm. *Economica*, 4(16):386–405, 1937.

Guido Matias Cortes, Nir Jaimovich, Christopher J. Nekarda, and Henry E. Siu. The micro and macro of disappearing routine jobs: A flows approach. *NBER Working Paper Series*, (20307):1–59, 2014.

Guido Matias Cortes, Nir Jaimovich, and Henry E. Siu. Disappearing routine jobs: Who, how, and why? *NBER Working Paper Series*, (22918):1–43, 2016.

John DiNardo, Nicole M. Fortin, and Thomas Lemieux. Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. *Econometrica*, 64(5):1001–1044, 1996.

Michael W.L. Elsby, Jennifer C. Smith, and Jonathan Wadsworth. The role of worker flows in the dynamics and distribution of uk unemployment. *Oxford Review of Economic Policy*, 27:338–363, 2011.

Michael W.L. Elsby, Bart Hobjin, and Aysegul Sahin. On the importance of the participation margin for labor market fluctuations. *Journal of Monetary Economics*, 72:64–82, 2015.

Harald Fadinger and Karin Mayr. Skill-biased technological change, unemployment, and brain drain. *Journal of the European Economic Association*, 12(2):397–431, 2015.

Gulia Faggio and Stephen Nickell. Inactivity among prime age men in the uk. *CEP Discussion Paper*, (673):1–21, 2005.

Bruce Fallick, John Haltiwanger, and Erika McEntarfer. Job-to-job flows and the consequences of job separations. *Finance and Economics Discussion Series*, (73): 1–43, 2012.

Christoper L. Foote and Richard W. Ryan. Labor market polarization over the business cycle. *Working Papers, Federal Reserve Bank of Boston*, (14-16):1–74, 2014.

Luca Foster, John Haltiwanger, and C.J. Krizan. Aggregate productivity growth: Lessons from microeconomic evidence. In Charles R. Hulten, Edwin R. Dean, and Michael J. Harper, editors, *New Developments in Productivity Analysis*, NBER Book Series Studies in Income and Wealth, pages 303–72. University of Chicago Press, 2001.

Lucai Foster, John Haltiwanger, and Chad Syverson. Reallocation, firm turnover, and efficiency: Selection on productivity and profitability? *American Economic Review*, 98(1):394–425, 2008.

Richard B. Freeman and Lawrence F. Katz. Introduction and summary. In Richard B. Freeman and Lawrence F. Katz, editors, *Differences and Changes in Wage Structures*, pages 1–22. University Chicago Press, 1995.

Shigeru Fujita and Guiseppe Moscarini. Recall and unemployment. *NBER Working Paper*, (19640):1–56, 2013.

Shigeru Fujita and Garey Ramey. Reassessing the shimer facts. *Federal Reserve Bank of Philadelphia Working Papers*, (07-2):1–20, 2007.

Shigeru Fujita and Garey Ramey. The cyclicality of separation and job finding rates. *International Economic Review*, 50(2):415–430, 2009.

Pedro Gomes. Labour market flows: Facts from the united kingdom. *Labour Economics*, 19:165–175, 2012.

Maarten Goos and Alan Manning. Lousy and lovely jobs: The rising job polarization of work in britain. *The Review of Economics and Statistics*, 89(1):118–133, 2007.

Maarten Goos, Alan Manning, and Anna Salomons. Job polarization in europe. *American Economic Review*, 99(2):58–63, 2009.

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

Peter Gottschalk and Mary Joyce. Cross-national differences in the rise in earnings inequality: Market and institutional factors. *The Review of Economics and Statistics*, 130(4):489–502, 1998.

Alfred Greiner, Jens Rubart, and Willi Semmler. Economic growth, skill-biased technical change and wage inequality: A model and estimations for the us and europe. *Journal of Macroeconomics*, 26(4):597–621, 2004.

Zvi Griliches. *RD and productivity: the econometric evidence*. University of Chicago Press, 1998.

Peter A. Hall and David Soskice. An introduction to varieties of capitalism. In Peter A. Hall and David Soskice, editors, *Varieties of Capitalism: The Institutional Foundations of Comparative Advantage*, chapter 1, pages 1–68. Oxford University Press, 2001.

Robert E. Hall and Charles I. Jones. Why do some countries produce so much more output per worker than others? *The Quarterly Journal of Economics*, 114(1): 83–116, 1999.

John Haltiwanger. Firm dynamics and productivity growth. *European Investment Bank Papers*, 16(1):116–136, 2011.

Craig Holmes and Ken Mayhew. The changing shape of the uk job market and its implications for the bottom half of earners. *Resolution Foundation*, pages 1–31, 2012.

Jingjing Huo. *How Nations Innovate: The Political Economy of Technological Innovations in Affluent Capitalist Economies*. Oxford University Press, 2015.

ifo Institute. Database for institutional comparisons in europe, 2013.

Nir Jaimovich and Henry E. Siu. The trend is the cycle: Job polarization and jobless recoveries. *NBER Working Paper Series*, (18334):1–53, 2012.

Stephen P. Jenkins. Survival analysis. Mimeo, Department of Economics, University of Essex, 2005.

Michal Jerzmanowski. Total factor productivity differences: Appropriate technology vs. efficiency. *European Economic Review*, 51(8):2080–2210, 2007.

Lawrence F. Katz and David H. Autor. Changes in the wage structure and earnings inequality. In Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, volume 3, pages 1464–1555. Elsevier, 1999.

Lawrence F. Katz and Kevin M. Murphy. Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics*, 107(1):35–78, 1992.

Lawrence F. Katz, Gary W. Loveman, and David G. Blanchflower. A comparison of changes in the structure of wages in four oecd countries. In Richard B. Freeman and Lawrence F. Katz, editors, *Differences and Changes in Wage Structures*, pages 25–66. University Chicago Press, 1995.

Nicholas M. Kiefer. Economic duration data and hazard functions. *Journal of Economic Literature*, 26(2):646–679, 1988.

Alan B. Krueger. How computers have changes the wage structure: Evidence from microdata, 1984-1989. *The Quarterly Journal of Economics*, 108(1):33–60, 1993.

Paul Krugman. Past and prospective causes of high unemployment. *Economic Review*, 79(4):23–43, 1994.

Stephen J. Nickelland Richard Layard. Labor market institutions and economic performance. In Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, volume 3, pages 3029–3084. Elsevier, 1999.

Thomas Lemieux. The changing nature of wage inequality. *NBER Working Paper*, (13523):1–40, 2007.

Frank Levy and Richard J. Murnane. U.s. earnings levels and earnings inequality: A review of recent trends and proposed explanations. *Journal of Economic Literature*, 30(3):1333–1381, 1992.

Stephn Machin and John Van Reenen. Technology and changes in skill structure: Evidence from seven oecd countries. *The Quarterly Journal of Economics*, 113 (4):1215–1244, 1998.

Guy Michaels, Ashwini Natraj, and John Van Reenen. Has ict polarized skill demand? evidence from eleven countries over 25 years. *Review of Economics and Statistics*, 96(1):60–77, 2014.

Joel Mokyr. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press, 1990.

Guiseppe Moscarini and Fabien Postel-Vinay. The contribution of large and small employers to job creation in times of high and low unemployment. *American Economic Review*, 102(6):2509–39, 2012.

Kevin M. Murphy, W. Craig Riddell, and Paul M. Romer. Wages, skills, and technology in the united states and canada. *NBER Working Paper Series*, (6638): 1–42, 1998.

Stephen Nickell, Luca Nunziata, and Wolfgang Ochel. Unemployment in the oecd sinde the 1960s. what do we know? *The Economic Journal*, 1:1–27, 2005.

OECD, 2015. URL `http://dx.doi.org/10.1787/data-00285-en`.

Mary O'Mahony, Catherine Robinson, and Michela Vecchi. The impact of ict on the demand for skilled labour: A cross-country comparison. *Labour Economics*, 15:1435–1450, 2007.

ONS. Annual survey of hours and earnings, 1997-2017: Secure access. [data collection], 2017a.

ONS. Emp01 sa: Full-time, part-time and temporary workers (seasonally adjusted), 2017b. URL `https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/employmentandemploy`

ONS. Time series: Gross domestic product: chained volume measures: Seasonally adjusted £m, 2017c. URL `https://www.ons.gov.uk/economy/grossdomesticproductgdp/timeseries/abmi/qna`.

ONS. New earnings survey panel dataset, 1975-2016: Secure access. [data collection], 2017d.

Joern-Steffen Pischke. Labor market institutions, wages, and investment: Review and implications. *CESifo Economic Studies*, 51(1):47–75, 2005.

Christopher A. Pissarides. *Equilibrium Unemployment Theory*. The MIT Press, 2 edition, 2000.

Thomas Pluemper and Vera E. Troeger. Efficient estimation of time-invariant and rarely changing variables in finite sample. *Political Analysis*, (15):124–139, 2007.

Andrea Salvatori. The anatomy of job polarisation in the uk. *IZA Discussion Paper Series*, (9193):1–53, 2015.

Wiemer Salverda and Daniele Checchi. Labour-market institutions and the dispersion of wage earnings. *IZA Discussion Paper Series*, (8220):1–140, 2014.

Joseph Alois Schumpeter. *Capitalism, Socialism and Democracy*, chapter The Process of Creative Destruction, pages 81–86. W. W. Norton Company, 2 edition, 2011.

Robert Shimer. Reassessing the ins and outs of unemployment. *Review of Economic Dynamics*, 15:127–148, 2012.

Christopher L. Smith. The dynamics of labor market polarization. *Finance and Economics Discussion Series. Board of Governors of the Federal Reserve System*, (2013-57):1–35, 2013.

Richard Upward and Peter Wright. Snakes or ladders? skill upgrading and occupational mobility in the us and the uk during the 1990s. *Research Paper Series Globalisation and Labour Markets. University of Nottingham*, (2007/38):1–37, 2007.

J. Visser. Ictwss data base, 2015. Version 5.0.

Oliver E. Williamson. *The Economic Institutions of Capitalism.* The Free Press, 1985.