The London School of Economics and Political Science

Essays in Labour Economics

Tiernan Evans

A thesis submitted to the Department of Economics of the London School of Economics and Political Science for the degree of Doctor of Philosophy

July 2025

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of approximately 30,000 words.

Statement of co-authored work

I confirm that Chapter 3 was jointly co-authored with Per-Anders Edin, Georg Graetz, Sofia Hernnäs, Guy Michaels, and I contributed 20% of this work. This paper was published in The Economic Journal in 2023.

Edin, Per-Anders, Tiernan Evans, Georg Graetz, Sofia Hernnäs, and Guy Michaels. "Individual consequences of occupational decline." The Economic Journal 133, no. 654 (2023): 2178-2209.

Statement of inclusion of previous work

I confirm that Chapter 2 of this thesis is a revised and extended version of the paper I submitted at the end of my Master of Research degree at the LSE in 2020.

Abstract

This thesis examines factors that influence economic and labour market outcomes throughout the life-cycle, with a focus on the importance of place.

The first chapter examines the role of local labour market conditions on educational take-up and human capital investment in England. I match administrative data on firms and students to document the variation in the level and field of skills demanded and study whether local skill demand shapes local students' educational attainment. I document a positive cross-sectional correlation between the skills demanded in local jobs and education choices. But, using a dynamic difference-in-difference strategy, I find at most a very muted response to large increases in local demand for degrees or specific skills for subsequent cohorts of students making educational investment decisions.

The second chapter turns to the question of geographic mobility, using PSID data to shed light on the generational dynamics of (internal) migration. I find that children born in a different state than at least one of their parents have approximately 20-25 percent higher short-term and medium-term interstate migration rates, and 50 percent higher lifetime mobility rates. These differences are robust to controlling for a wide set of observables, and are consistent across subgroups by education and gender. These findings have significant implications for spatial sorting models and our understanding of intergenerational transmission of economic opportunity.

In the third chapter, we assess the career earnings losses that individual Swedish workers suffered when their occupations' employment declined. Our estimates show that occupational decline reduced mean cumulative earnings from 1986–2013 by no more than 2%–5%, with larger losses for those initially at the bottom of their occupations' earnings distributions. This loss reflects a combination of reduced earnings conditional on employment, reduced years of employment and increased time spent in unemployment and retraining.

Acknowledgments

Doing a PhD is a difficult process, and it is thanks to the support of many people that this thesis was possible. This thesis has benefited greatly from the support of my colleagues, friends, and family.

Firstly, I want to thank my supervisor, Guy Michaels, whose mentorship has been instrumental throughout this journey. His generous sharing of time, experience, and insights has been invaluable in my development as a researcher and in refining my projects. I am immensely grateful for his patience, support, and encouragement, which have defined this experience. I also extend my sincere appreciation to my advisor, Camille Landais, whose astute observations have significantly shaped my research.

I've had the pleasure of working with many excellent researchers outside of LSE in the course of my PhD. I am grateful to my coauthors who brought me into a project (that now features in this thesis) which allowed me to get my start in research and first foray into the publication process with an excellent team. I also learned a great deal from my work with Camille Terrier, Parag Pathak, and Kevin Ren. Camille also offered invaluable support in navigating the final stages of the PhD.

I thank Xavier Jaravel, Steve Machin, Alan Manning, Sandra McNally, Henry Overman, Steve Pischke, Kate Smith, Johannes Spinnewijn and Alwyn Young for their valuable feedback and suggestions that improved this thesis. I also gratefully acknowledge the insights provided in the CEP Education and departmental seminars.

My colleagues at LSE and the CEP were also a definitive part of my PhD experience. I am grateful to Ihsaan Bassier, Shania Bhalotia, Alexandre Desbuquois, Gaia Dossi, Arnoud Dyevre, Andres Felipe Fajardo Ramirez, Shadi Farahzadi, Thomas Glinnan, Nilmini Herath, Gustave Kennedi, Peter Lambert, Vivian Liu, Thomas Monk, Kamila Nowakowicz, Tanner Regan, Rebecca Rose, Guglielmo Ventura, Maria Ventura, and many other LSE and visiting researchers whose friendship and academic exchange have made these years truly special.

My friends in London outside of the PhD and academic life have been a crucial part of this journey. I owe much to Doug for his unwavering support and to Matt for the stimulating intellectual conversations that reminded me of the joy in academic pursuit. The softball community has been a welcome respite from PhD life, and I am grateful for the friends I made there, especially Alex—for lending an ear whenever I needed and encouraging sensible choices—and Lorraine—for ensuring I maintained a social life and for her constant companionship and support.

My time in London and as a PhD student has not been without its challenges, and my friend Louise Bernard has been there every step of the way. Her fierce intelligence and unwavering support have been a beacon of light this journey, and I am fortunate to have had the chance to learn from her. I am also thankful for my friendship with Cecila Wood. From our eerily similar backgrounds to the haggis-stuffed Thanksgiving turkey that has become an annual tradition, I could always count on her for great company and a thoughtful perspective.

Finally, my deepest gratitude extends to my family for encouraging me to pursue any path in life, even if it took me across an ocean. My parents, Martha and Charles, my grandmother Gretchen, and my siblings, Garrett and Kiri, have been an enormous source of love and support.

Table of Contents

List of Tables

List of Figures

Chapter 1

Local labour markets and skill acquisition

1.1 Introduction

There are large and persistent differences in educational attainment and skills across locations within countries. These spatial inequalities have recently drawn increased attention for their potential economic and political impacts, as well as for the individual consequences for those in disadvantaged areas. However, more insight is needed as to the ways in which various local characteristics are perpetuated. One important decision facing young people is the choice of education - both in terms of the type and level of education they pursue and the field in which they study.

Within England there are large, persistent, differences in educational attainment across areas, which have garnered significant concern and media attention. The gaps between London and rural areas have been documented at least as far back as the mid-nineteenth-century (White 2007) and have persisted into the twenty-first century (Overman and Xu 2022). Despite England's relatively small size, location is sticky: 65% of university graduates and 85% of non-graduates live in the same Travel to Work Area (TTWA) in which they grew up (Britton, Waltmann, Xu, et al. 2021). These characteristics lead to a natural research question: How responsive are students' educational investment decisions to local skill demand in the labour market in which they grow up?

While I study this question in the context of England (for reasons discussed below), this setting is not unusual in its spatial disparities. For example, in the US, the share of adults with university degrees ranges from less than 27% in Louisiana and Mississippi to more than 42% in Massachusetts and Vermont, and household income in the poorest states is 55% of that in the richest (U.S. Census Bureau 2019). And despite overall higher (though declining) geographic mobility, 60% of Americans live in the same state in which they were born Jia, Molloy, Smith, and Wozniak 2023.

Governments have aimed to address spatial inequalities in part through place-based industrial policy, using both government jobs and incentives for private investment. In the UK, these policies have included the Places for Growth scheme, enacted, in part, to "[ensure] economic growth and job opportunities are more evenly distributed across the UK" by moving more than 20,000 Civil Service jobs outside of London (Department for Business 2021). In the US, the Helping Infrastructure Restore the Economy (HIRE) Act, proposed in 2019, would have required the relocation of several federal government agencies away from Washington, D.C.. More recently, the Creating Helpful Incentives to Produce Semiconductors (CHIPS) and Science Act of 2022 provided $39 billion of funding to support local hubs and foster the growth of the semiconductor industry.

Understanding whether such policies can spur local skill accumulation is important. The most direct route by which demand may affect locals' skills is through cohorts of students who may change their skill acquisition in response. In the case where the elasticity of education choices to the jobs available is small, these policies may instead primarily change local skill composition through (costly) moves of workers from other locations.

To better understand how students' educational investment decisions respond to local skill demand in the labour market in which they grow up, I turn to England. In addition to the spatial disparities described above, England has many advantageous institutional features for this question. Firstly, school funding is determined nationally on a per-pupil basis, meaning an expansion in local employment or an increase in house prices does not increase education funding. This stands in sharp contrast to other contexts where local taxes provide the bulk of education funding, confounding estimates of student response[1]. Furthermore, both academic and vocational qualifications fit into a standardized framework from before the end of compulsory schooling, allowing me to compare more specific skills at an earlier stage of education. The comparability of qualifications across time and space, as well as comprehensive data on student results, allows me to examine the field of study, student achievement, and level of education. Finally, administrative data on students and firms is available to study.

In this paper, I use rich administrative data on both establishments and students in England to evaluate the education response of students to changes in the composition of local skill demand. I examine the role that existing differences in skill demand play in local students' educational attainment: their results; level of education; and fields of study. I then exploit large changes in local demand for skills - either qualifications in a specific field or at the degree level - to measure the short- to medium-run response of local students. These changes come from changing local industry composition; I show that, following a large change, local demand for that skill remains at a significantly higher level than it was in the preceding years.

To understand local skill demand, I use the Business Structure Database (BSD), an administrative dataset covering an estimated 99% of economic activity in the UK. This establishment-level data provides information on the precise location, industry and employment of each establishment. I combine this with survey data on the qualifications held by employees in each sector to measure the mix of skills (by level and field), used in each labour market. The granular nature of the data allows me to consider other geographic levels, as well as changes coming from specific events.

On the student side, I use English administrative data on the universe of students who attended state (publicly funded) schools. This data, which is linked to university records, allows me to observe the full educational path - including the subjects studied and results - of students from primary through tertiary education.

My primary unit of analysis is the TTWA—these regions are defined by the ONS using census data such that the bulk (at least 75%, or 66.7% in larger areas) of their resident population work within the same area. These regions are mutually exhaustive and collectively exclusive, covering the whole of the UK. Conceptually, they are similar to Commuting Zones and can be thought of as local labour markets[2]. Thus, the type of employment present in the TTWA can be thought of as characterising the labour market which students in an area will face. In some specifications, I also consider results at the Postcode District (PD) level. This smaller unit is roughly equivalent to a zip code; it has an average area of 25 square miles.

In the first part of the paper, I document the spatial differences in industries and their associated

1. Consider an influx of well-paid jobs requiring degrees; the resulting increase in local property tax base would mechanically increase school funding, an important contributor to school quality, which is associated with higher levels of educational attainment (Eckert and Kleineberg 2021).

2. While recent work (e.g. **mp2017<empty citation>**, Monte, Redding, and Rossi-Hansberg (2018)) emphasises the importance of local spatial linkages, this would be better incorporated using LEO data to empirically understand the expected labour market for these cohorts

skills and descriptively relate them to the spatial differences in education attainment. In this section, I show that students' education choices, in both level and field, are strongly correlated with the estimated skill distribution of local employment. I then move to a staggered event study design exploiting large changes in local skill demand to identify causal estimates of students' response.

First, I look at how the estimated share of local jobs requiring a degree relates to student enrolment. I find that, in the cross-sectional correlations, compared to areas in the 25th percentile of share of degree jobs, those in the 75th have 10pp more students enrolling in A-levels (compared to a national average of 45%) and 8pp more students enrolling in university (compared to a national average of 27%). I then consider the choice of field of study in university; here, for every additional 10pp of local employment in a field (compared to other TTWA), there is a 3pp difference in university enrolment in that field.

I then exploit large changes in local area skill composition to understand whether and when students respond to changing local conditions. I consider 1-year increases of more than 1pp in local degree employment in the TTWA where a student began secondary school; these changes are persistent. Despite the strong correlational relationships, I find no significant changes in enrollment in A-Levels or University following an increase in degree employment, nor change in field enrolment following an expansion of local employment in that field. To bound the response, I turn to students' immediate surroundings and consider the same change at the postcode group level. Here I also find no evidence for an increase in A-level enrolment, and can bound the response from above by a 2.5pp increase in enrolment for every 1pp increase in the share of jobs requiring a degree. The response for university enrolment is similarly small with the upper bound at 1.8pp for the same increase; this translates to just 5 additional students enrolled in university for every 100 new jobs requiring degrees.

In response to local increases in field share of more than 0.75% of local employment, I first show that these expansions are persistent - the level of employment in that field in the subsequent years remains significantly above its initial level. I then examine the student response. At the TTWA level, I find no evidence of increased field enrolment. Again turning to the postcode group level, I can bound the response from above at a 1pp increase in student enrolment for every 1.5pp increase in field employment.

Given the strong correlation between the skills used in local jobs and students' education choices, the small short-run response of students to changing circumstances warrants further consideration. It could be the case that the period of adjustment is longer than I observe in my data; for example, if local teachers know more about the pathways that previous cohorts of students took, or if students beliefs about the labour market they will enter update slowly. It could also be the case that students do not view local changes as relevant for their later job prospects.

The muted response of local students has important implications for policies targeting regional inequalities. The magnitude of responses I find are not economically meaningful: the response is not enough to fill newly-created high-skill jobs with local students. If that is the goal of these policies, additional intervention may be needed to shape education choices. However, an important caveat is that my results relate only to academic qualifications; it is also relevant to understand whether other skills, such as those requiring vocational training, which are more likely to be used locally are likely to respond to changes.

This also has implications for the impact of these policies on local growth. If local students do not fill the demand, the vacancies created by programs intended to move high-skill jobs will need to be filled otherwise. People may move from areas with a high supply of in-demand skills; however, these moves are costly, and empirically, internal migration rates are not large enough to fill the vacancies created by large increases in specific industries. These jobs may instead be filled by locals with a different (sub-optimal) skill mix, reducing productivity of the industry. An important avenue for future research is understanding how these skill gaps are bridged.

This paper contributes to understanding in several areas of research. First, it provides an empirical basis for the nascent literature endogenizing education choice in models of spatial inequality and mobility. Prior to Eckert and Kleineberg (2021), which to my knowledge is the first to include this adjustment margin in their model, this literature took skill supply as a fixed local characteristic. In this paper, I exploit local changes to establish a supply elasticity for skills in a context where the supply of education does not vary in response to local income changes.

Secondly, this paper relates to an emerging literature understanding how students' choices respond to local labour market conditions. Most closely related is contemporaneous work examining the impact of STEM jobs on high school and college course choice and major in the US (Mather and Smith 2024). In contrast to this paper, I focus on isolating the student response in a setting where local conditions do not directly impact secondary education provision. Other papers in this area (e.g. Conzelmann et al. (2023) and Weinstein (2022)) focus on college major choice, using university location to determine the relevant labour market. I consider instead students' earlier location and include their decision whether to attend university in my outcomes.

This work also expands on prior work examining the impact of specific local shocks on educational attainment in various contexts. The existing literature finds that education is a substitute for labour in response to single-sector shocks caused by natural resource booms (Emery, Ferrer, and Green (2012), Maurer (2019), and Kovalenko (2023)) and trade policy (Atkin 2016). I differ from this literature by considering the changes in local composition coming from all sectors, rather than a single industry or policy, and by considering more nuanced educational choices in addition to continued enrolment. My results are also consistent with prior work finding a (permanent) dip in the attainment of those who graduate into a housing-induced local labour market boom (Betts and McFarland (1995), Charles, Hurst, and Notowidigdo (2018)).

Also related is the literature on the elasticity of education choice to expected earnings. This literature has focused primarily on college major choice in the US, finding a small but significant elasticity[3]. However, this literature does not explicitly consider the cost of relocation or the spatial distribution of jobs which is the focus of my work.

The findings here also relate to recent work on place-based policies, the hazard cost of which depends on the spatial concentration and willingness to move of the intended transfer recipients (Gaubert, Kline, and Yagan 2021). My research points to an additional margin —education choice— in which people in a given location may be entrenched and provides evidence that industrial policy may not

3. Altonji, Arcidiacono, and Maurel (2016) provides a good review of the findings.

change these outcomes for locals.

My results also bring results from agglomeration literature into a new light. Recent work has highlighted the importance of spillovers from significant local industrial investment on the local and total cost of investment policies (Kline and Moretti 2014). This work finds the largest (positive) spillovers from investment in areas which share worker flows with the industry being funded (i.e. already have a pool of qualified workers) (Greenstone, Hornbeck, and Moretti 2010). The muted response I find from local students sheds light on why having an existing local pool of qualified workers is important.

The rest of this paper is organised as follows. Section 1.2 describes the institutional context of English education. Section 1.3 describes the datasets and measures used in the analysis, and presents descriptive statistics. Section 1.4 describes the empirical strategy and assumptions required for the identification strategy. Section 1.5 presents the findings for attainment and field choice. Section 1.6 describes additional specifications and extensions. Section 1.7 concludes and highlights potential avenues for future research.

1.2 Institutional context

In the UK, education is specialised, and qualifications are standardized at a young age. Education is regulated at the national level; that is, each of England, Scotland, Wales and Northern Ireland have their own requirements. In this paper, I study students in England.

Figure 1.1 shows the typical progression of students through the English education system. The first opportunity for subject choice comes in Key Stage 4, when students are ages 14-16. At this phase of their education, students choose which subjects to study for their General Certificate of Secondary Education (GCSE) examinations. These exams, typically taken at the end of Year 11 (ages 15-16), are nationally standardised and offered in a variety of subjects. All students are required to take exams in English, mathematics, and science, and schools are required to provide at least one offering in each of arts, humanities, business, and language. Thus, these choices are comparable for all students in England.

Schooling in England is compulsory until age 16. Education is compulsory to age 18; this was increased from 16 to 17 in 2013, and from 17 to 18 in 2015. Post-16 qualifications are standardised across the UK and regulated through the Regulated Qualifications Framework which covers both academic and vocational qualifications. In England, the relevant framework is the National Qualification Framework (NQF). During this period, students have several alternatives for education: academic-track A-levels; technical education; vocational training; on-the-job training; or a mix of volunteering and part-time employment.

In this paper, I focus on whether students continue on the academic track. In contrast to vocational programmes, such as apprenticeships, which require local employment in the field to be viable, the supply of academic qualifications is less dependent on local demand. Academic disparities, including A-level results and university share, have also been the focus of recent public and political attention.

Following the period of compulsory education, students can choose to continue their education either academically through university courses (which generally require completion of A-levels) or through continuing vocational education. Like academic qualifications, vocational qualifications are regulated and standardized, and thus comparable across time, level, and field.

Students can select up to 5 subjects to study in A-levels but typically take 3 exams, the number required for most university courses. University departments post descriptions of the typical profile of subjects and grades required for admission, which informs students' choices about which subjects to take and where to apply. In practice, this means that the pathways students can take can be restricted by their choice of A-level subjects[4].

In England, school funding is determined on a per-pupil basis according to the National Funding Formula (NFF); this stands in sharp contrast to the US, where school funding comes primarily from the local tax base. While this formula is adjusted for local economic conditions[5], it is unrelated to local tax revenue. The funding is distributed to the Local Authority, which then allocates it to schools according to its own formula.

This policy results in school funding on a per-pupil basis being more equitable across schools and less responsive to local conditions than in other contexts. In England, during this time, the 90/10 ratio of per-pupil funding at the school level ranged from 1.3 to 1.6, with the most deprived local authorities having the highest level of funding (Belfield and Sibieta 2016). This stands in sharp contrast to the US where differences are much larger[6] and positively correlated with local economic conditions, shutting down an important channel of inter-generational human capital transmission. Conceptually, this allows me to measure the change in student attainment holding the local education supply fixed.

1.3 Data and descriptive statistics

1.3.1 Student data

Student information comes from administrative data from the Department for Education (DfE). Data on qualifications attained during compulsory education, as well as demographic and location information, comes from the National Pupil Database (NPD) (Department For Education 2023). This is an administrative panel dataset on the census of students attending state-funded secondary schools in England beginning in the 2001/02 school year. Data on university attainment comes from the Higher Education (HE) dataset, which can be linked to records from the NPD. This student data is collected for all students studying a qualification above level 3 (A-Levels or equivalent vocational qualification) at a reporting HE provider.

I attribute to each student the local labour market conditions in the TTWA (PD) in which they lived when they began their GCSE studies (age 14-15, the beginning of Key Stage 4). Throughout, I will refer

4. Although schools are not required to offer a specific selection of A-levels, students generally have access to a broad variety of subjects. The subjects typically required for most university programmes are offered broadly throughout the country. Figures 1.A2 through 1.A8 show A-level subjects which commonly appear in university requirements. Each map shows the student-weighted share of schools in each TTWA which offer the degree (TTWA with fewer than 3 A-level providers are omitted).

5. The funding level is adjusted for local (teacher) hiring costs, a measure of income deprivation in students, and the level of remoteness from other services.

6. For comparison, the 90/10 ratio of per pupil funding in the US for elementary and secondary schools at the more aggregated school district level was 2.4 in 2016 (U.S. Census Bureau 2021).

to this as the TTWA (or PD) which the student is from[7].

Figure 1.2a shows the average share of students from each TTWA who enrol in A-levels. The share of students enrolling in A-levels ranges from less than 50% in some coastal areas to nearly 70% in parts of the southeast. This pattern is mirrored in Figure 1.2b, showing university enrolment, with the share enrolled varying from less than 25% to nearly 50% along the same geographic lines.

In addition to enrolment and attainment, these datasets provide specific information about each subject taken (A-levels) and degree program (university). I map each of these to the fields coded in the Labour Force Survey (LFS). The mapping process is described in more detail in section 1.C.1. Figure 1.3 shows the field of university enrolment that is most disproportionately popular in each TTWA; that is, the field for which the share of college enrolment is most above the national average. This map shows the geographic clustering of different fields, with Business degrees being overrepresented in London and the surrounding areas and broadly-applicable Arts and Humanities degrees overrepresented throughout the southeast.

1.3.2 Business data

Data on local establishments comes from the Business Structure Database (BSD) (Office for National Statistics 2024), which is an administrative dataset covering all businesses liable for VAT and/or with at least one member of staff registered for the PAYE tax collection system. This dataset contains approximately 2 million observations annually, covering an estimated 99% of UK economic activity. I observe the location, industry, and employment count at each establishment, as well as the firm to which it belongs.

Although this is an administrative dataset, there are still missing values for many entities. Because of the identification strategy employed, these have the potential to create multiple false 'events' in a given location. To alleviate this problem, I linearly interpolate these values using the surrounding years. This is described in more detail in the data appendix.

1.3.3 Other data

To link the employment from the BSD to the skills mix in the DfE data, I use the existing skill composition of young workers (ages 25-34) in each industry from the Labour Force Survey (LFS) (Office For National Statistics 2023). The LFS is a nationally representative survey designed to provide information on the UK labour market. In addition to employment information, including detailed industry, it includes detailed information on qualifications held by individuals. I use information on the level and field of the highest level of education. The sample of younger workers is used to more closely approximate the skills requirements for the current students, and for increased consistency in qualification standards over time.

7. Student mobility during this phase of their schooling is generally low. In each cohort, more than 98% of students observed in year 12 live in the same TTWA they did in year 10.

1.3.4 Data linking

Because the BSD includes only headcount by employer, not information about employees, I must impute the local qualification distribution using data from the LFS. I describe this process in more detail in section 1.4.1. Linking the BSD to the LFS by industry allows me to translate the industry employment information in the BSD into estimated employment by qualification.

I aggregate the student data by geography, calculating the total number of students who take each type of qualification discussed above. I then match this geography x time panel to the business data to create a novel dataset linking student outcomes and business demand.

1.4 Empirical strategy

To estimate the student response to changes in local labour market conditions, I consider students living in location (local labour market) l at time t, measure changes skill s in location l, and observe the students' academic outcomes in later years.

1.4.1 Measuring local skill demand

The first step in measuring local skill demand is determining the qualifications used in each industry. I compute this using the empirical distribution of held by workers ages 25-34 in industry j in the LFS. Specifically, I sum the weighted person count by qualification type (field or level) for each industry and divide by the total weighted person count in the industry. This creates a skill share vector, $skills_j$, for each industry j.

I also compute the local industry employment shares, $emp_{j,l,t}$, using BSD data. This is simply the total employment in industry j and location l at time t.

I then combine these elements to aggregate these industry skill vectors and calculate the skill demand in location l at time t. This is calculated as the employment-weighted sum of skills vectors in the industries in that location and time.

$$skills_{l,t} = \sum_{j \in J} emp_{j,l,t} * skills_j \tag{1.1}$$

Here, $skills_{l,t}$ is the vector of skills used in employment in location l at time t. $emp_{j,l,t}$ is the employment in industry j in l at time t. $skills_j$ is industry skill vector discussed above. The assumption embedded in this measurement is that the skill requirements for new jobs in the industry are the same as those held by existing young workers.

The industry-specific skills vectors are derived from LFS data as described in Section 1.3.3. Thus, the actual local skill composition is estimated using the national industry-average skill mix. While this approach was a necessity due to data availability, it has the advantage that it abstracts one level from skill supply decisions in firm moves. That is, the local area need not have an existing supply of workers with a specific skill for that skill to be recognised as in-demand by the industry. This approach can be thought of as analogous to a shift-share instrument, where there shares are the industry skill mix and the shifters are the local industry composition. This interpretation is discussed in more detail in section 1.5.2.

For tractability, I consider separately the level and field of skills demanded. For outcomes relating to student enrolment decisions and attainment, I consider only the level of qualifications required. For these specifications, the skills vector is $[Degree, Non-degree]$. For specifications relating to field choice, I use the field composition of qualifications[8] seen in Figure 1.4b.

Figure 1.4a shows the average expected share of employment in a TTWA requiring a degree. As with the student data, rates of university attendance are much higher in the South East. Figure 1.4b shows the field of study that is most disproportionately represented in the industry mix in each TTWA. As with the student data, Business degrees are overrepresented in London. Areas with lower degree employment also see a disproportionately high share of activities typically associated with rural areas, as well as services (health and education) that are provided locally.

Average skill share

In my correlational results, I consider the average skill share in the three years prior to a decision. I take this measure to smooth any transient changes in labour market shares and because the exact timing of the relevant labour market characteristics is ambiguous.

Changes in the overall local skill composition

I utilize the change in local skill mix coming from changing local industrial composition. The change I compute is the one-year difference in skill share. $skills_{s,l,t}$ is the component of the skills vector $skills_{l,t}$ in skill s, and $skills_{l,t}$ is size of the workforce in location l at time t. Thus, $\Delta skills_{s,l,t}$ is the one year change in the percentage of workers with a given skill.

$$\Delta skills_{s,l,t} = \frac{skills_{s,l,t}}{||skills_{l,t}||} - \frac{skills_{s,l,t-1}}{||skills_{l,t-1}||}$$

To better understand the dynamics of changes in skill composition, I construct an indicator based on the change in skills in the local area. This allows me to compare cohorts who are yet to make a decision with those already past the start of the qualification.

$$D_{s,l,t} \equiv \mathbb{1}(\Delta skills_{s,l,t} > \Theta)$$

Change threshold Θ	**TTWA**	**PD**
Degree share	0.010	0.015
Field share	0.050	0.075

I chose a value of Θ to maximize the number of locations which experienced exactly one shock. The results and interpretation are not sensitive to alternative threshold choices.

8. The LFS attributes field only to qualifications level 3 (A-Levels or equivalent vocational qualification) and above in the NQF, thus, the field mix I use is that of qualifications at this level or higher.

To test the validity and observe the skill share dynamics of this change, I consider the following specification:

$$skills_{s,l,t} = \sum_{\tau=-4}^{6} \left(\beta_{\tau}^{emp} D_{s,l,t+\tau} \right) + \gamma_{s,t} + \alpha_{s,l} + \varepsilon_{s,l,t}$$

Figure 1.5 shows the dynamics of degree share employment around a large increase in the degree share. Figure 1.6 shows the same for fields of study. While, in both cases, the share in the years preceding the increase is significantly above that in the year immediately prior, the years after the increase are significantly higher than the earlier level. The overall change shows only a modest attenuation in the 5 subsequent years and remains well above the pre-increase levels. This indicates a sustained change in the share of employment requiring the specified skill.

1.4.2 Student response

First, for the correlational estimates, I consider how students' choices relate to the level of employment in a skill in their TTWA in the 3 years preceding the decision point, $skills_{s',l,t+\tau}$.

$$Sh_{s,l,t} = \beta \sum_{\tau=-3}^{-1} \left(skills_{s',l,t+\tau} \right) / 3 + \gamma_{s,t} + \varepsilon_{s,l,t} \tag{1.2}$$

$Sh_{s,l,t}$ is the share of student cohort in location l at time t who choose to study for qualification s. τ is 0 when the student is in the first year of schooling after the decision is made. For example, for the decision to attend university, $\tau = 0$ in a student's last year of A-levels. Take-up of a skill may change in response to demand changes for s, or for related skills s'; for example, A-level enrolment (s) is a prerequisite for degree enrolment, so an increase in jobs requiring degrees (s') can be used on the right-hand side. In all specifications, the data is at the Location x Skill x Year level. Observations are weighted by the total number of students in the TTWA in the initial year (2002). Standard errors are clustered at the TTWA level to account for the correlation of field shares within a location at a given time, and the correlation across time within a location from the averaging of local employment shares.

For the causal estimates, I estimate the following Difference-in-Differences specification:

$$Sh_{s,l,t} = \sum_{\tau=-3}^{4} \left(\beta_{\tau}^{incr} D_{s',l,t+\tau} \right) + \gamma_{s,t} + \alpha_{s,l} + \varepsilon_{s,l,t} \tag{1.3}$$

Here, the right-hand side variables of interest are lags and leads of the dummy for large increases described above. I also include Skill x Time and TTWA x Skill fixed effects. The unit of observation, weighting, and standard errors are the same as in the correlational estimates; here it is necessary to cluster across time within a location due to the local skill share fixed effect ($\alpha_{s,l}$).

1.5 Results

In this section, I first show the relationship between local skill demand and student education outcomes. I then discuss the assumptions required for the causal estimates that follow and present students' responses to changes in their local labour market.

1.5.1 Correlational evidence

Attainment and local degree share

First, I consider the relationship between local degree employment and student attainment. Table 1.1 shows the correlations at the TTWA level in both levels and differences[9]. TTWAs with a higher share of degree employment have a significantly higher share of students enroling in A-levels, more A-level passes (per student from that area), and a higher share of students enroling in university. Compared to TTWAs in the 25th percentile, those in the 75th have 10pp more students enrolling in A-levels and 8pp more students enroling in university. Looking at the 5-year change in local degree employment, the sign of the coefficients remains positive, but the magnitude is much smaller and they are no longer significant. Most of the differences come from variations between, rather than within, TTWAs.

Field choice and local field share

I next consider the relationship between local field employment and students' field choice. Table 1.2 shows the correlations at the TTWA level in both levels and differences. TTWAs with a higher share of field employment have a significantly higher share of students enroling in that field in university; the relationship does not hold for A-levels, which are much broader subjects. For every 10pp higher the field share of an area is compared to others, the share of students from that area studying in that field is 3pp higher. Looking at the 5-year change in local field employment, the sign of the coefficients reverses, indicating that students may substitute education for employment within a field.

1.5.2 Inference and estimation challenges

To understand the dynamics of the changing labour demand, I use the estimators developed in De Chaisemartin and D'haultfœuille (2023). For these estimates to be interpreted as causal, there must be no anticipation effects, and the parallel trends assumption must hold.

In my specifications, $t = 0$ is the cohort that had just made a decision about the outcome variable, and thus, in the absence of anticipation effects, is unaffected. I report coefficients for the three preceding cohorts (who were already past the decision threshold at the time of treatment) as well and see estimates indistinguishable from 0 in all cases. To the extent the immediately preceding cohort anticipates the employment change (perhaps due to newspaper articles, etc.), the earlier cohorts would have to have the same anticipation, which seems unlikely given the nature of the changes.

In my specifications, I allow for nonparametric differences in the evolution of different fields over time. The parallel trends assumption is therefore within skill (e.g. the expected evolution of Natural Science degrees is the same for treated and untreated groups; it need not be the same as the evolution of Arts degrees). One potential cause for violation of the parallel trends assumption would occur if firm activity driving the events was a response to changing local circumstances. For example, if firms saw an increasing share of students studying for relevant skills, and expanded as a result; triggering an event in areas that were different from untreated areas. However, this would require firms to have detailed

9. Table 1.A1, shows how these margins respond to local labour market conditions as a whole.

knowledge of student choices at a local level, and an assumption that this increase would persist. I also do not observe any evidence for this scenario in the pre-event periods.

The variation I exploit can also be interpreted as a shift-share instrument for local skill demand, where the shares are the national skill vectors in each industry and the shifter is the change in local industrial composition. This interpretation helps to alleviate concerns about the endogeneity of increases in a specific skill, as the measured change is mediated by the industrial skill composition. However, by using this approach, actual changes in local skill demand are not fully captured. If, for example, the national (or regional) skills used in an industry are not representative of the skills required to be held by workers, the variation I use would not be relevant to students preparing to enter the labour force.

The other plausible shift-share instrument for this work is to shift the local industry composition (and thus skill composition) using the change in national industry composition (i.e. assuming any national shifts are proportional across space). This would address endogeneity concerns about the growth of industries in specific locations as opposed to others. However, it has the drawback (particularly for contractions in local demand for a skill) that it entangles local and national demand.

1.5.3 Causal evidence

I find muted responses to local labour market conditions, and responses to hyper-local (postcode district) conditions different than those at the local labour market (TTWA) level.

Attainment and local degree share

I first consider how enrolment and attainment respond to an increase in the share of local jobs requiring degrees. A-levels are the most common path to university enrolment; therefore, if the choice to enrol in A-levels was primarily driven by forward-looking skill concerns, enrolment would grow in response to increased demand for jobs requiring degrees.

Figure 1.7 shows how enrolment in A-levels changes in response to an increase in local degree employment, as specified in equation 1.3. I find no evidence that students' choice to enrol in A-levels responds to an increase local degree demand.

I then turn to the composition in students' more immediate vicinity. Figure 1.8, shows the response at the PD level. The positive point estimates increase a change of around 0.35pp in the first two years and 0.90pp in the following two. The confidence intervals rule out an increase of more than 2.5pp, equivalent to one student per cohort for every 10 new jobs. With an average increase in local employment of 2.5pp in the share of jobs requiring a degree, this rules out an increase even at replacement level for the new skill composition.

The next margin of choice students face is whether to attend university. Figure 1.9 shows how the share of students who enrol in a degree program changes in response to increase in local degree employment. Again, I find no evidence that enrolment changes following an increase in degree employment [10]
.

10. Figure 1.A1, shows how these margins respond to local labour market conditions as a whole. I find no evidence that students substitute away from university enrolment when the local economy expands.

Field choice and local field share

Finally, I consider the field choice of university students. Figure 1.10 shows the share of students who study that subject in university. Here, I also find no evidence that, in the short term, students respond to changing local demand. The estimates also bound the change at no more than an 0.4pp increase for every 1pp increase in field share. At the PD level, there is a borderline-significant increase in field enrolment, as seen in figure 1.11. However, the magnitude is small: it is bounded from above by a 0.25pp increase in enrolment for every 1pp increase in field share in jobs. This amounts to 1 new student in a field for every 200 new jobs.

All together, my results show that any short-run student response to changing local conditions is muted, and that the change in sills from local students is insufficient to meet the new demand.

1.5.4 Heterogeneity

I now consider how these results differ for different types of students. First, I consider the type of school where the students attended A-levels. In England, different types of schools have different levels of autonomy. Here I consider separately students enrolled in community schools, which are funded and fully administered by the local authority, and those enrolled in academies (including converters) which are funded by the DfE but independent of the Local Authority. Figure 1.12 shows the response limiting the sample to students who, in Year 12, were enrolled in an academy. Here, I find noisy estimates and cannot rule out no response. However, the magnitude of the point estimates is much larger than for the total sample, and warrants further study. Figure 1.13 shows the results for only students enrolled in community schools. As in the full sample, the magnitude o the response is small enough to be indistinguishable from 0 - I find no evidence that students in these schools respond to changes in local demand.

I then consider whether these results differ by the type of university attended. This is an endogenous choice for students, so the causal interpretation is less clear. I consider the subject choices of students who attend Russell Group universities and those who attend former polytechnics. Those who attend Russell Group universities are much more likely to live outside their home TTWA after university (Britton, Waltmann, Xu, et al. 2021), and therefore may be expected to respond less to local demand. Figure 1.14 shows the field choice response for students enrolled in Russell Group universities, and 1.15 shows the same for those enrolled in former polytechnics. In both cases, there is no evidence for an increase in field enrolment following an expansion of the local field share.

1.6 Additional specifications

1.6.1 Different skills

My current analysis focuses on academic qualifications — A-levels and university — but there is also a unified qualification framework for vocational and technical skills. Students who do not attend university are more likely to remain in the area where they grow up, so the local conditions may be more relevant for them. These skills are relevant for policies such as Reshore UK, intended to relocate manufacturing jobs from abroad to the UK. The response to demand for these skills may be different due to different training modes, and may also be more dependent on local training supply. The endogenity of the availability of these programmes will require careful empirical analysis.

1.6.2 Student characteristics

In addition to using different variations in the skills demanded, it is also relevant to examine how different types of students respond to changes in local demand.

Earlier performance

Graduates from the most selective group of universities[11] are more mobile than graduates of other universities, with those from the least selective universities being the least likely to move (Britton, Waltmann, Xu, et al. 2021). Another fruitful extension could therefore be to investigate how these responses vary with students' earlier performance by splitting students into groups based on their results before the decision to enter A-levels.

Geographic mobility

In the current analysis, I consider all students that are living in a place at a given time (before the local economic changes). I can also restrict the analysis to students who remain in that area. Conversely, I could identify the area effects on students who move with a measure of 'exposure' to different subjects à la Chetty, Hendren, and Katz (2016).

Demographics

While the primary focus of this analysis is the drivers of inequality across space, the drivers of education choice and responsiveness to local changes may also contribute to inequality within a location. Students who are more responsive may differentially benefit from higher returns to in-demand local skills (Altonji, Blom, and Meghir 2012 Early et al. 2020). Although I do not observe the parents' industry, I will use an indicator for whether the student was eligible for Free School Meals (FSM) as a measure of familial background. This will allow me to identify whether the responsiveness is different by parental background.

1.6.3 Broader changes

This analysis has focused on local changes, controlling for national circumstances, rather than on broader demand. Particularly for university decisions, the overall prospects may have an important impact. Another avenue for future research is using a shift-share instrument to how national shocks, such as import competition, affect students' choices.

11. The most selective universities are defined as the Russell Group, a self-selected group of twenty four research universities. These universities receive the bulk of all research funding.

1.7 Conclusion

In this paper, I study spatial differences in student attainment, achievement, and field choice. I find that students' choices are highly correlated with those used in jobs in their local area, but, in the short run, do not significantly respond to changes in local demand. This has important consequences for the results of policies implemented with the goal of bringing good jobs to left-behind communities.

Multiple mechanisms could be at play causing these differences: the local changes may not be salient, the types of skills required may not be those that students had considered, or students may be considering other factors in their choice of study. Regarding salience, I find that even the most local of changes (on average within 5 miles of students' location) do not significantly affect take-up. To further address this concern, future research can instead consider large expansions or contractions of prominent local firms, which may be more noticeable to local students.

While there is no mechanical relationship between one cohort and the next there, could also be inertia at the local or school level hampering changes. For example, teachers may be more familiar with the pathways that previous cohorts took. Furthermore, there are other local factors, such as parental composition, that may push students to follow the more traditional local path (Ventura 2023). Understanding these factors leading to local persistence of education choices is a potentially fruitful area for further research.

While the English setting is ideal for understanding student responses holding constant local education provision, this needs to be considered when applying of this research to other contexts. Expectations of the total student response to changing local circumstances would also need to consider local skill provision, as the response I measure is that of students when local provision does not change. This may be complicated by latent demand for skills which were not previously offered, resulting in in higher take-up of these avenues than I observe in England.

The muted response of local students to changing conditions leaves many questions open for future policy research. If the goal of these policies is for the local population to gain valuable skills, it is important to know whether coordinating job growth policies with targeted upskilling programs may induce a response not seen with changing job skill composition alone. It is also relevant to understand whether other skills, such as those requiring vocational training, which are more likely to be used locally are likely to respond to changes.

Bibliography for Chapter 1

Altonji, Joseph G, Peter Arcidiacono, and Arnaud Maurel. 2016. "The analysis of field choice in college and graduate school: Determinants and wage effects." In *Handbook of the Economics of Education,* 5:305–396. Elsevier.

Altonji, Joseph G, Erica Blom, and Costas Meghir. 2012. "Heterogeneity in human capital investments: High school curriculum, college major, and careers." *Annu. Rev. Econ.* 4 (1): 185–223.

Atkin, David. 2016. "Endogenous skill acquisition and export manufacturing in Mexico." *American Economic Review* 106 (8): 2046–2085.

Belfield, Chris, and Luke Sibieta. 2016. "Long-Run Trends in School Spending in England: IFS Report R115." *Institute for Fiscal Studies.*

Betts, Julian R, and Laurel L McFarland. 1995. "Safe port in a storm: The impact of labor market conditions on community college enrollments." *Journal of Human resources,* 741–765.

Britton, Jack, Ben Waltmann, Xiaowei Xu, et al. 2021. "London calling? Higher education, geographical mobility and early-career earnings: Research report: September 2021."

Charles, Kerwin Kofi, Erik Hurst, and Matthew J Notowidigdo. 2018. "Housing Booms and Busts, Labor Market Opportunities, and College Attendance." *American Economic Review* 108 (10): 2947–94.

Chetty, Raj, Nathaniel Hendren, and Lawrence F Katz. 2016. "The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment." *American Economic Review* 106 (4): 855–902.

Conzelmann, Johnathan G, Steven W Hemelt, Brad Hershbein, Shawn M Martin, Andrew Simon, and Kevin M Stange. 2023. *Skills, Majors, and Jobs: Does Higher Education Respond?* Technical report. National Bureau of Economic Research.

De Chaisemartin, Clement, and Xavier D'haultfœuille. 2023. "Two-way fixed effects and differences-in-differences estimators with several treatments." *Journal of Econometrics* 236 (2): 105480.

Department for Business, Energy Industrial Strategy. 2021. *Business department and Home Office to open up almost 3,000 Civil Service job roles outside of London.* https://www.gov.uk/government/news/business-department-and-home-office-to-open-up-almost-3000-civil-service-job-roles-outside-of-london. [Accessed 11-11-2024].

Department For Education. 2023. *National Pupil Database.*

Early, Erin, Sarah Miller, Laura Dunne, Allen Thurston, and Mehmet Filiz. 2020. "The influence of socio-economic background and gender on school attainment in the United Kingdom: A systematic review." *Review of Education* 8 (1): 120–152.

Eckert, Fabian, and Tatjana Kleineberg. 2021. "Saving the American Dream? Education Policies in Spatial General Equilibrium."

Emery, JC Herbert, Ana Ferrer, and David Green. 2012. "Long-term consequences of natural resource booms for human capital accumulation." *Ilr Review* 65 (3): 708–734.

Gaubert, Cecile, Patrick M Kline, and Danny Yagan. 2021. *Place-based redistribution.* Technical report. National Bureau of Economic Research.

Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. “Identifying agglomeration spillovers: Evidence from winners and losers of large plant openings.” *Journal of political economy* 118 (3): 536–598.

Jia, Ning, Raven Molloy, Christopher Smith, and Abigail Wozniak. 2023. “The Economics of Internal Migration: Advances and Policy Questions.” *Journal of Economic Literature* 61 (1): 144–80.

Kline, Patrick, and Enrico Moretti. 2014. “Local economic development, agglomeration economies, and the big push: 100 years of evidence from the Tennessee Valley Authority.” *The Quarterly journal of economics* 129 (1): 275–331.

Kovalenko, Alina. 2023. “Natural resource booms, human capital, and earnings: Evidence from linked education and employment records.” *American Economic Journal: Applied Economics* 15 (2): 184–217.

Mather, Ryan, and Andrew Smith. 2024. “Hometown Labor Markets and Degree Choice.”

Maurer, Stephan E. 2019. “Oil discoveries and education provision in the Postbellum South.” *Economics of Education Review* 73:101925.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg. 2018. “Commuting, Migration, and Local Employment Elasticities.” *American Economic Review* 108 (12): 3855–90.

Office For National Statistics. 2023. *Labour Force Survey.*

Office for National Statistics. 2024. *Business Structure Database - UK.*

Overman, Henry, and Xiaowei Xu. 2022. “Spatial disparities across labour markets.” *Institute for Fiscal Studies* 2.

U.S. Census Bureau. 2019. *S1501: Educational Attainment [American Community Survey 5-Year Estimates Subject Tables 2019].* https://data.census.gov/table/ACSST5Y2019.S1501?q=Education&y=2019.

———. 2021. *2016 Public Elementary-Secondary Education Finance Data.*

Ventura, Maria. 2023. *Following in the Family Footsteps: Incidence and Returns of Occupational Persistence.*

Weinstein, Russell. 2022. “Local labor markets and human capital investments.” *Journal of Human Resources* 57 (5): 1498–1525.

White, Jerry. 2007. *London in the Nineteenth Century:’a human awful wonder of God’*. Chap. VII Fleet Street: City of Words. Random House.

Appendix

Figure 1.1: England education timeline

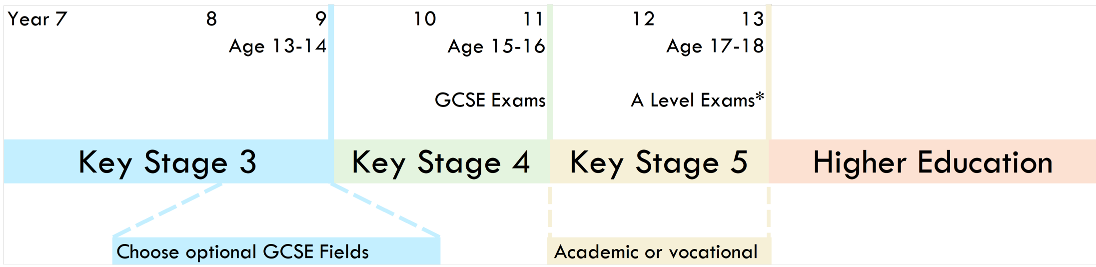

Figure 1.2: Local student enrolment

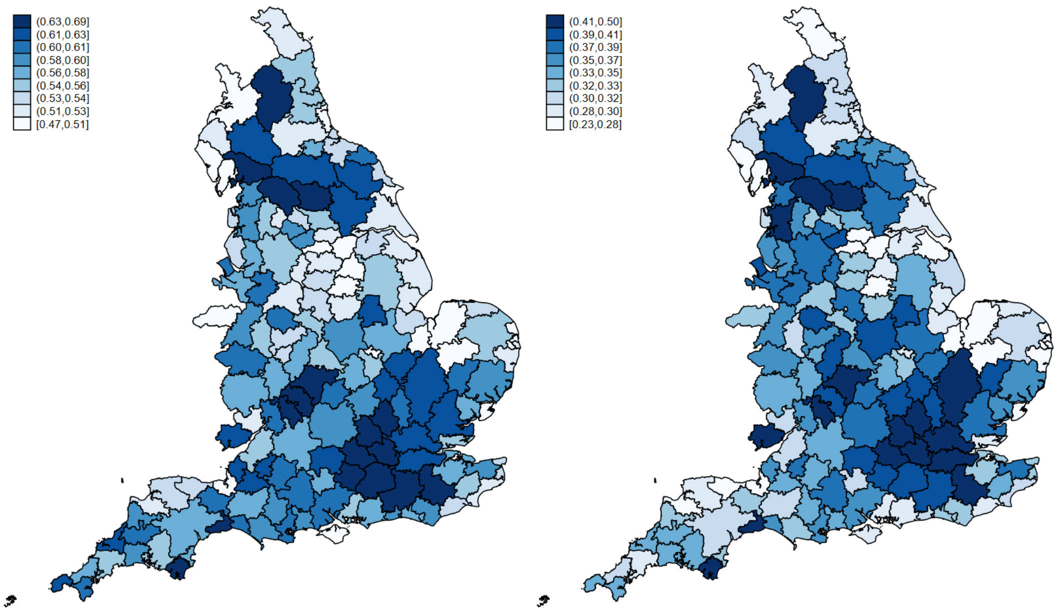

(a) A-level enrolment

(b) University enrolment

Student information from the NPD. Cohorts that were in Year 10 from 2004-2013 are included. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10.

Figure 1.3: Local students' most prominent university field choice

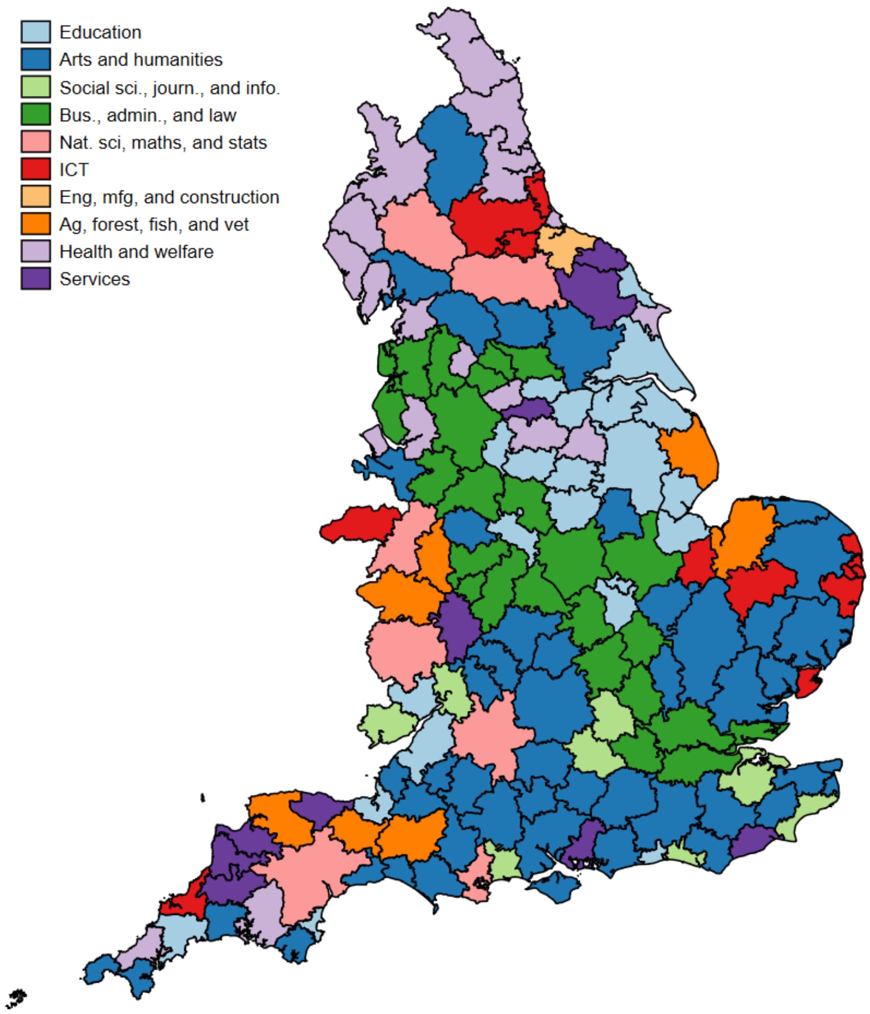

For each TTWA, the map shows the field of study that is most overrepresented compared to national levels. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10.

Figure 1.4: Local employment shares

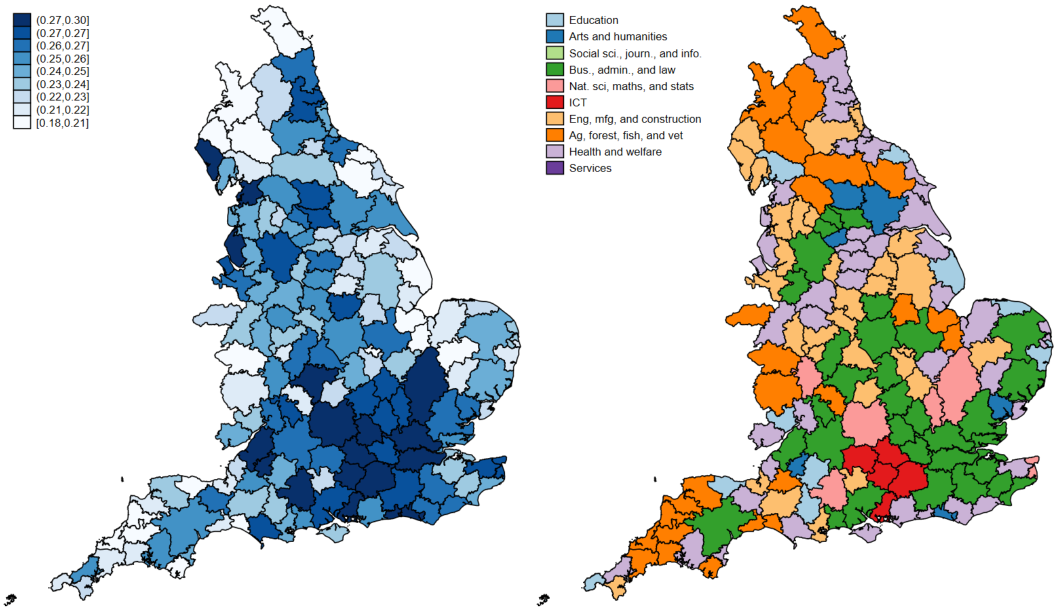

(a) Degree employment (b) Field most above national average

Employment shares from the BSD, combined with the LFS. For each TTWA, the map in panel (b) shows the field of employees in that location that is most overrepresented compared to national levels.

Figure 1.5: Degree employment shares increase following degree expansion

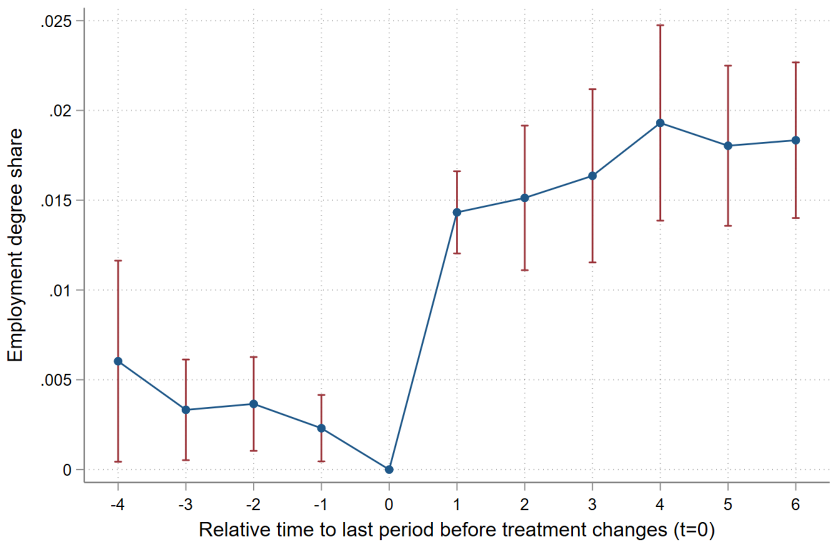

Data at the TTWA x Year level. Employment shares and changes from the BSD.

Figure 1.6: Field employment shares increase following field expansion

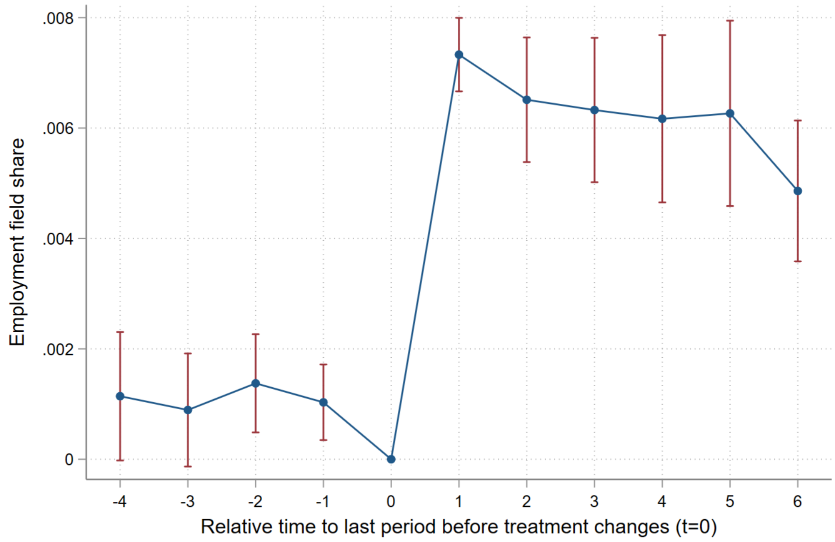

Data at the TTWA x Year x Qualification level. Employment shares and changes from the BSD.

Figure 1.7: A-level enrolment following an expansion of local (TTWA) degree employment

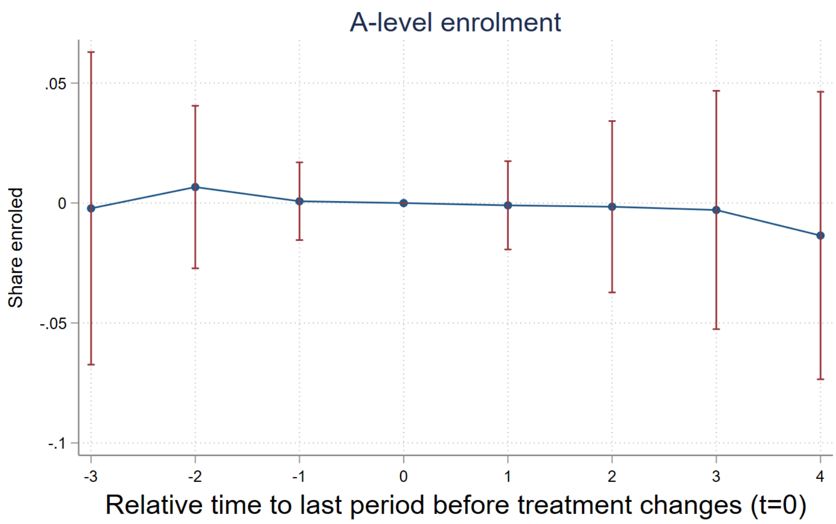

Data at the TTWA x Year level. Employment shares and changes from the BSD. Student information from the NPD. Cohorts that were in Year 10 from 2004-2013 are included. RHS variable used is an indicator for changes of more than 1.0pp in the share of jobs in an area that employ a worker with a degree. Standard errors are two-way clustered by base year TTWA.

Figure 1.8: A-level enrolment following an expansion of Postcode District degree employment

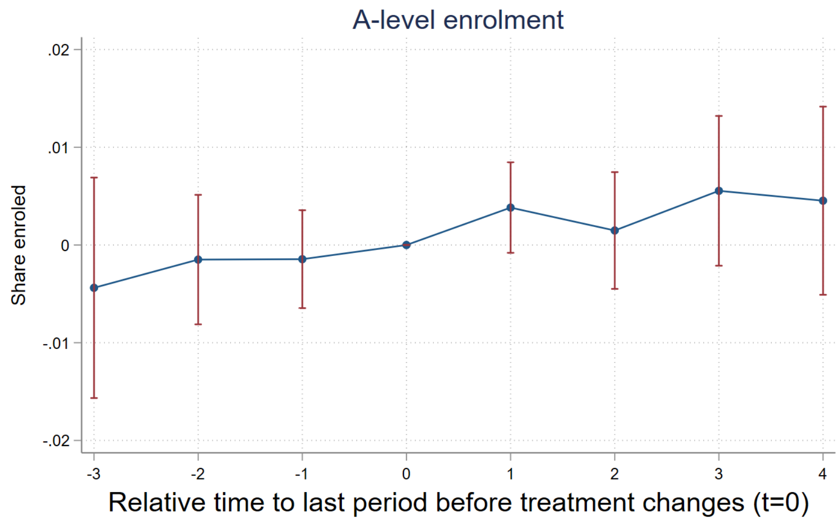

Data at the Postcode District x Year level. Employment shares and changes from the BSD, student shares from the NPD. Student information from the NPD. Cohorts that were in Year 10 from 2004-2013 are included. RHS variable used is an indicator for changes of more than 1.5pp in the share of jobs in an area that employ a worker with a degree. Standard errors are two-way clustered by base year TTWA.

Figure 1.9: Degree enrolment following an expansion of local degree employment

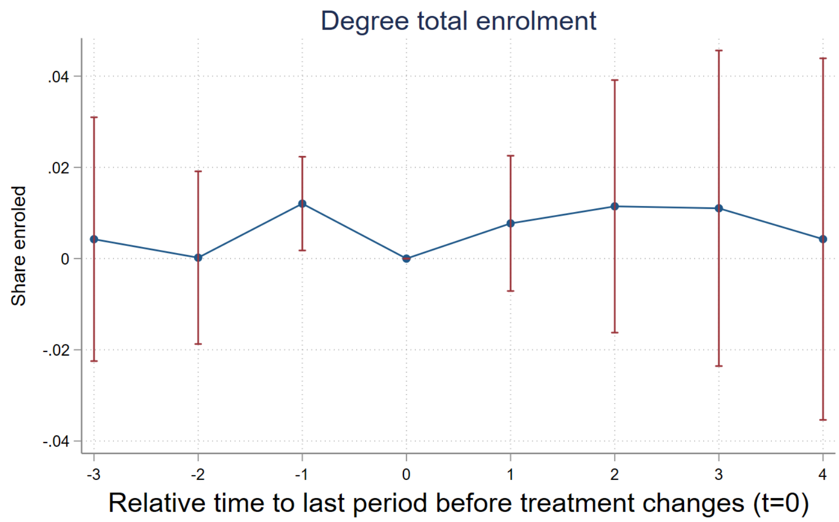

Data at the TTWA x Year level. Employment shares and changes from the BSD. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. RHS variable used is an indicator for changes of more than 1.0pp in the share of jobs in an area that employ a worker with a degree. Standard errors are two-way clustered by base year TTWA.

Figure 1.10: University field selection following an expansion of local field employment

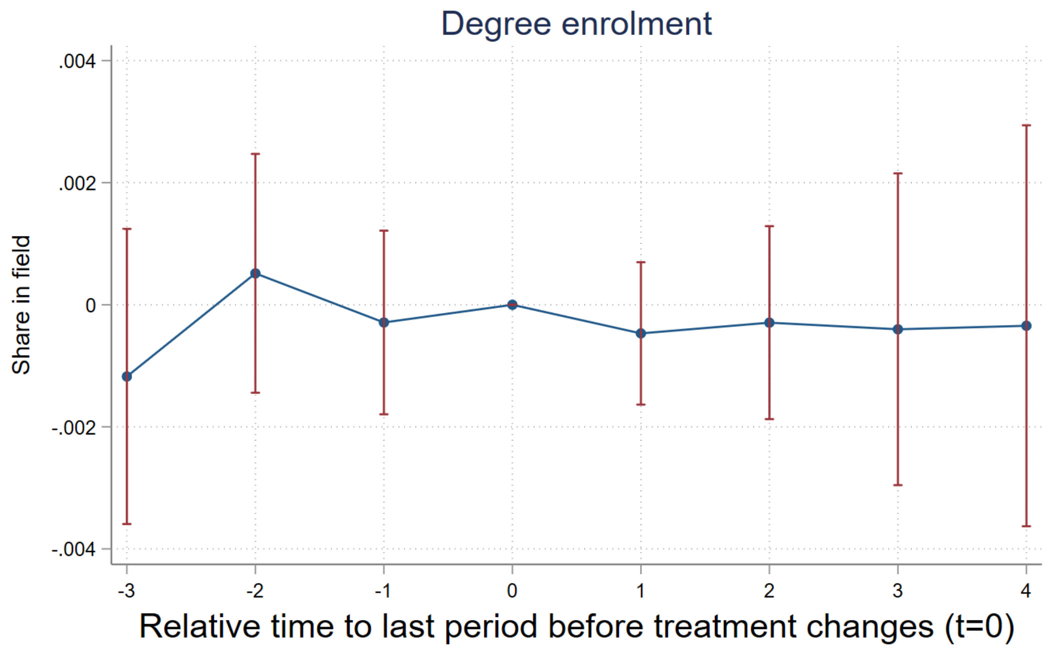

Data at the TTWA x Year x Qualification level. Employment shares and changes from the BSD. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. RHS variable used is an indicator for changes of more than 0.5pp in the share of jobs in an area that employ a worker with a qualification in the field. Standard errors are two-way clustered by base year TTWA.

Figure 1.11: University field selection following an expansion of Postcode District field employment

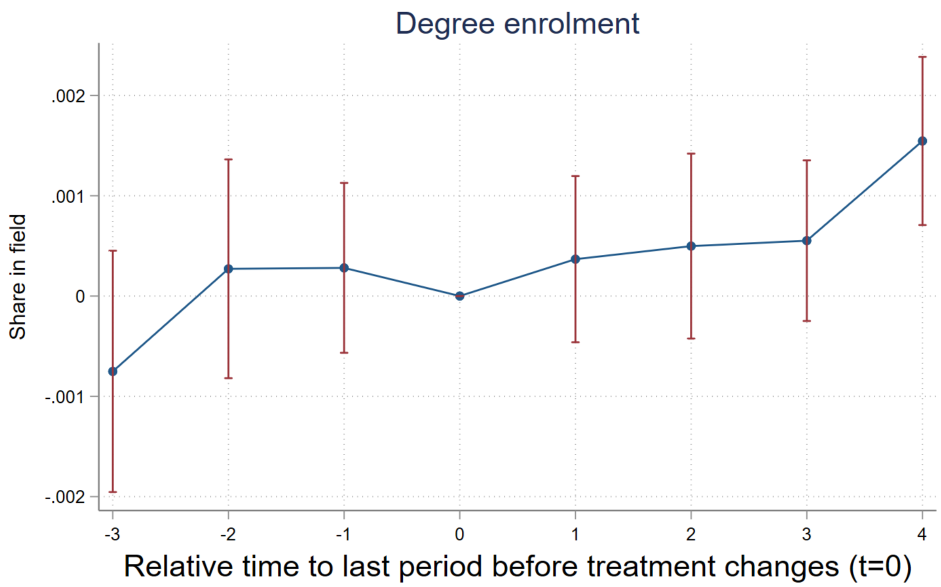

Data at the Postcode District x Year x Qualification level. Employment shares and changes from the BSD, student shares from the NPD. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. RHS variable used is an indicator for changes of more than 0.75pp in the share of jobs in an area that employ a worker with a qualification in the field. Standard errors are two-way clustered by base year TTWA.

Figure 1.12: University field selection for academy students following an expansion of local field employment

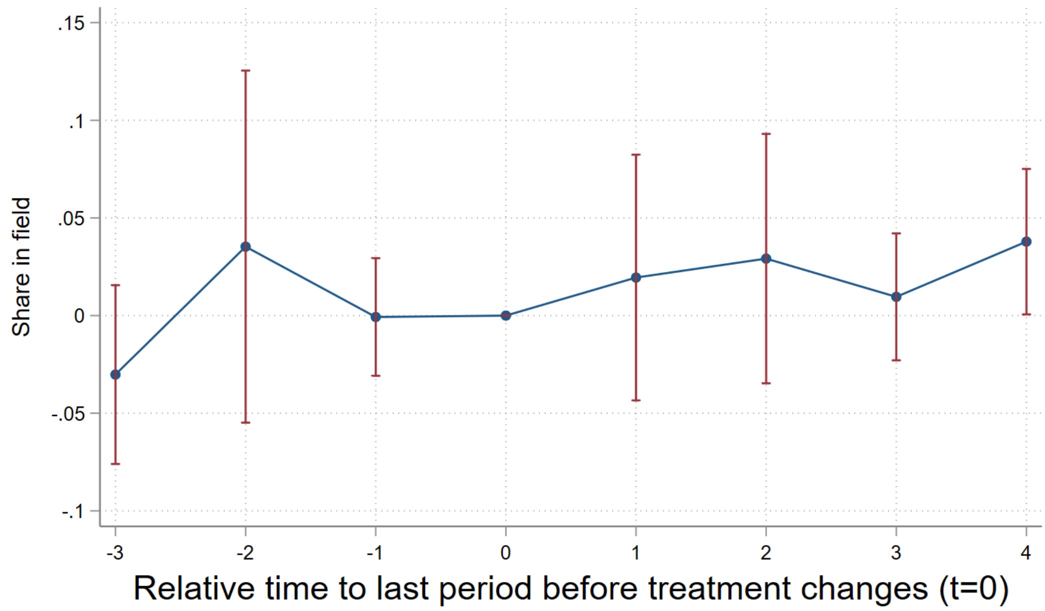

Data at the TTWA x Year x Qualification level. Employment shares and changes from the BSD. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. Only students who attended an academy in Year 12 are included. RHS variable used is an indicator for changes of more than 0.5pp in the share of jobs in an area that employ a worker with a qualification in the field. Standard errors are two-way clustered by base year TTWA.

Figure 1.13: University field selection for community school students following an expansion of local field employment

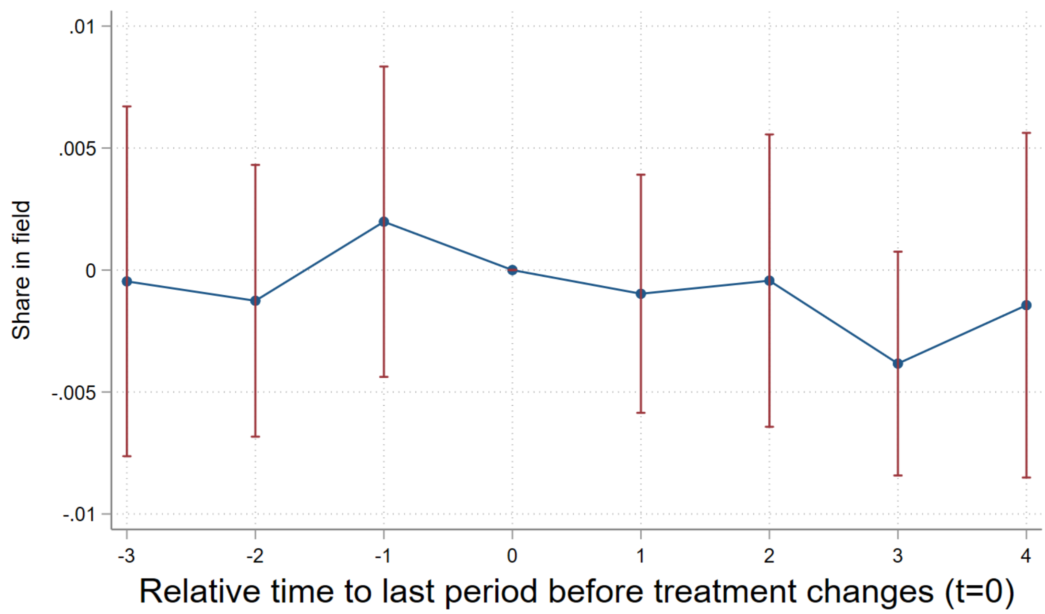

Data at the TTWA x Year x Qualification level. Employment shares and changes from the BSD. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. Only students who attended a community school in Year 12 are included. RHS variable used is an indicator for changes of more than 0.5pp in the share of jobs in an area that employ a worker with a qualification in the field. Standard errors are two-way clustered by base year TTWA.

Figure 1.14: University field selection for Russell Group students following an expansion of local field employment

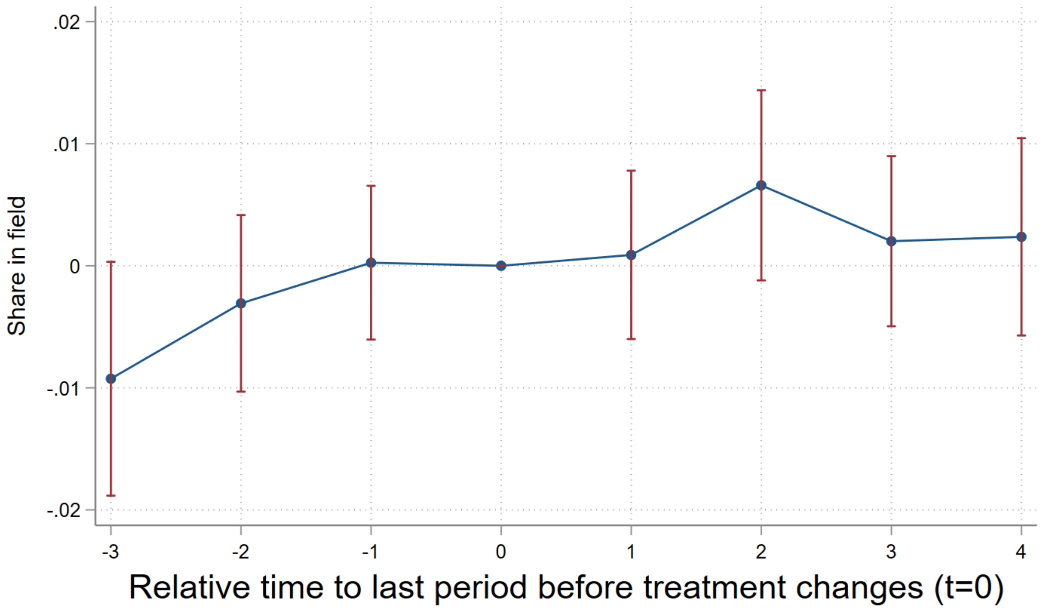

Data at the TTWA x Year x Qualification level. Employment shares and changes from the BSD. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. Only students who attended a Russell Group are included. RHS variable used is an indicator for changes of more than 0.5pp in the share of jobs in an area that employ a worker with a qualification in the field. Standard errors are two-way clustered by base year TTWA.

Figure 1.15: University field selection for students at former polytechnic universities following an expansion of local field employment

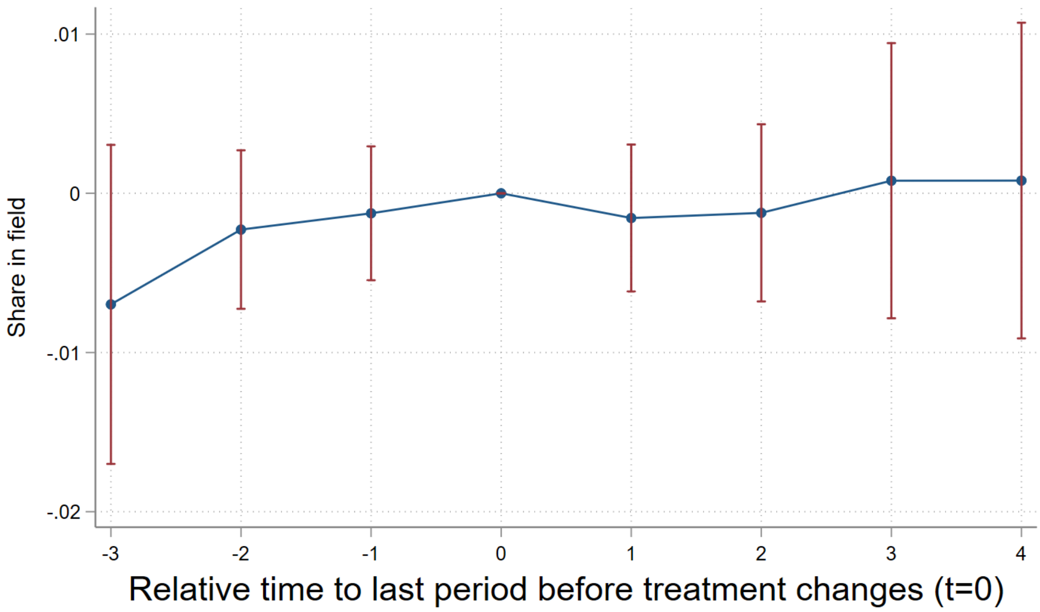

Data at the TTWA x Year x Qualification level. Employment shares and changes from the BSD. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. Only students who attended a polytechnic university are included. RHS variable used is an indicator for changes of more than 0.5pp in the share of jobs in an area that employ a worker with a qualification in the field. Standard errors are two-way clustered by base year TTWA.

Appendix 1.A Results

Figure 1.A1: Degree enrolment, total economic conditions

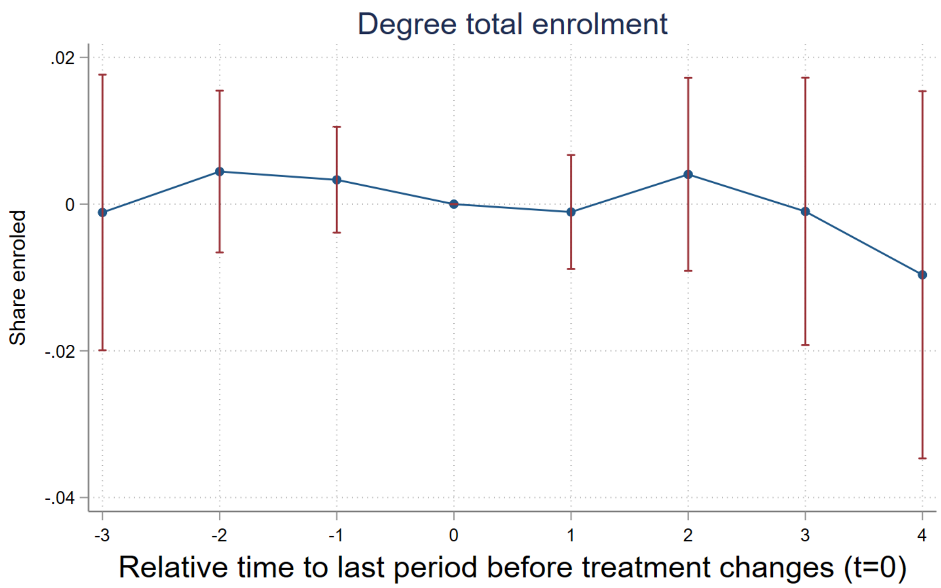

Data at the TTWA x Year level. Employment shares and changes from the BSD. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. Standard errors are two-way clustered by base year TTWA.

Table 1.1: Relationship between local degree share and student enrolment

	A-Level Enrol.	A-Level Passes	Uni Enrol.	A-Level Enrol.	A-Level passes	Uni Enrol.
Share degree emp.	0.544***	1.566**	0.409**			
	(0.118)	(0.549)	(0.125)			
Δ_5 Share degree emp.				0.157	0.390	0.0890
				(0.154)	(0.660)	(0.157)
Year FE	Y	Y	Y	Y	Y	Y
N	1705	1705	1705	930	930	930

Data at the TTWA x Year level. Employment shares and changes from the BSD. Student information from the NPD. Cohorts that were in Year 10 from 2004-2013 are included. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. Results are weighted by the number of students in Year 10 in the TTWA in 2002. Standard errors are two-way clustered by base year TTWA.

Table 1.2: Relationship between local field share and student subject choice

	A-Level field	Uni Field	A-Level Field	Uni Field
Share field emp.	0.107	0.316***		
	(0.149)	(0.0390)		
Δ_5 Share field emp.			-0.523*	-0.0835
			(0.238)	(0.111)
Field x Year FE	Y	Y	Y	Y
N	15345	18600	8370	10850

Data at the TTWA x Year x Qualification level. Employment shares and changes from the BSD. Student information from the NPD. Cohorts that were in Year 10 from 2004-2013 are included. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10. Results are weighted by the number of students in Year 10 in the TTWA in 2002. Standard errors are two-way clustered by base year TTWA.

Figure 1.A2: Student-weighted share of schools offering A-level biology

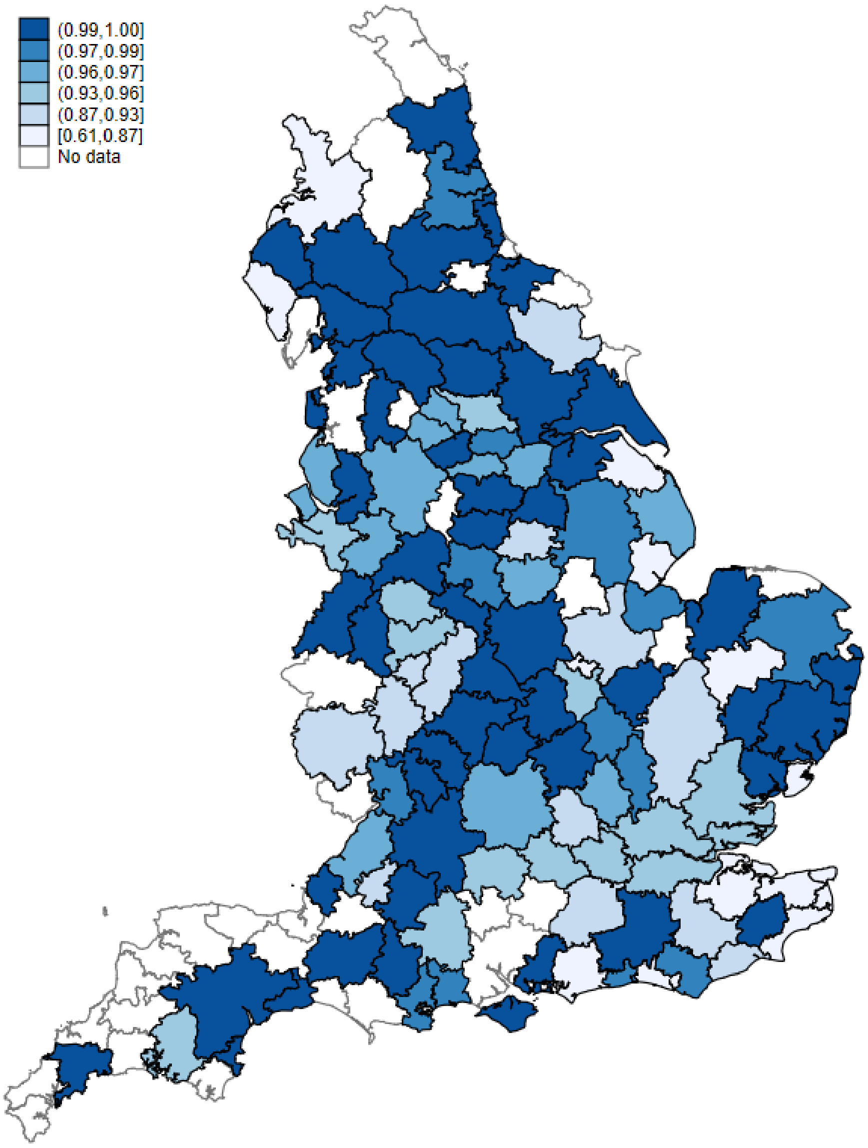

For each TTWA, the map shows the student-weighted share of schools which, in 2009 offered A-level biology.

Figure 1.A3: Student-weighted share of schools offering A-level chemistry

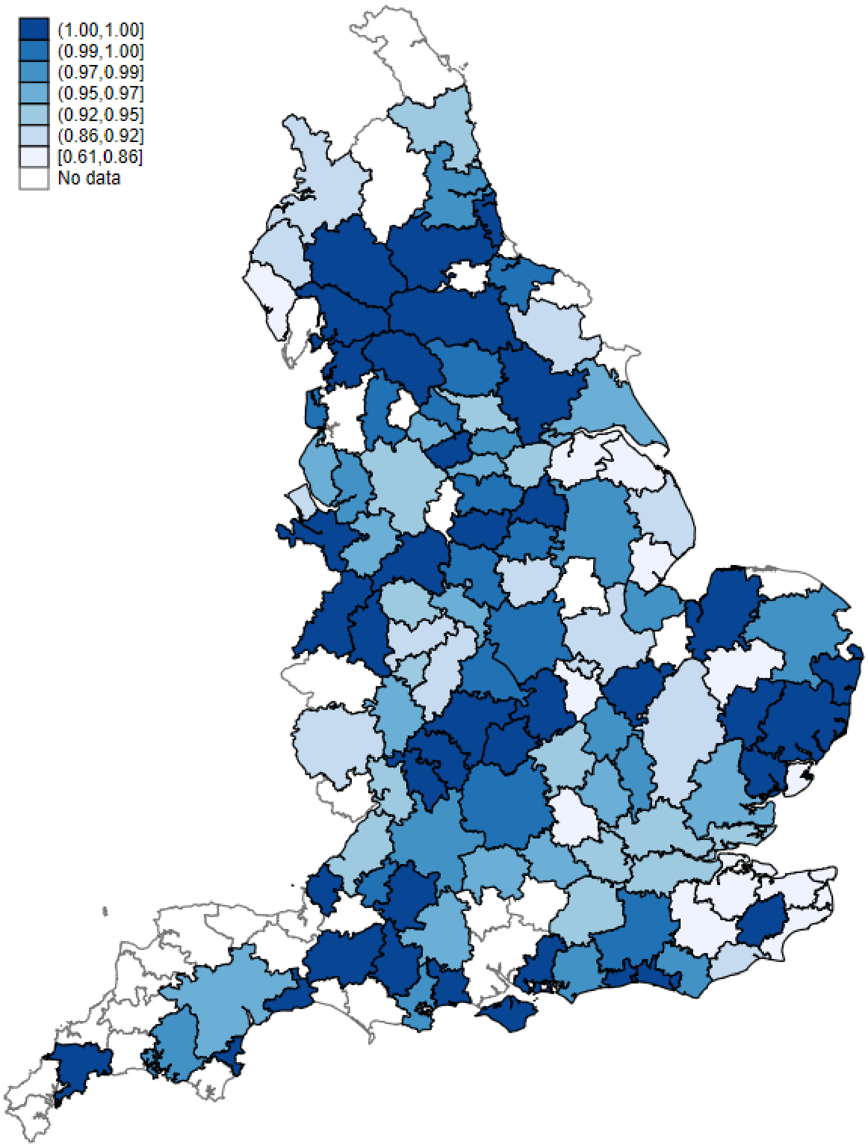

For each TTWA, the map shows the student-weighted share of schools which, in 2009 offered A-level chemistry.

Figure 1.A4: Student-weighted share of schools offering A-level english literature

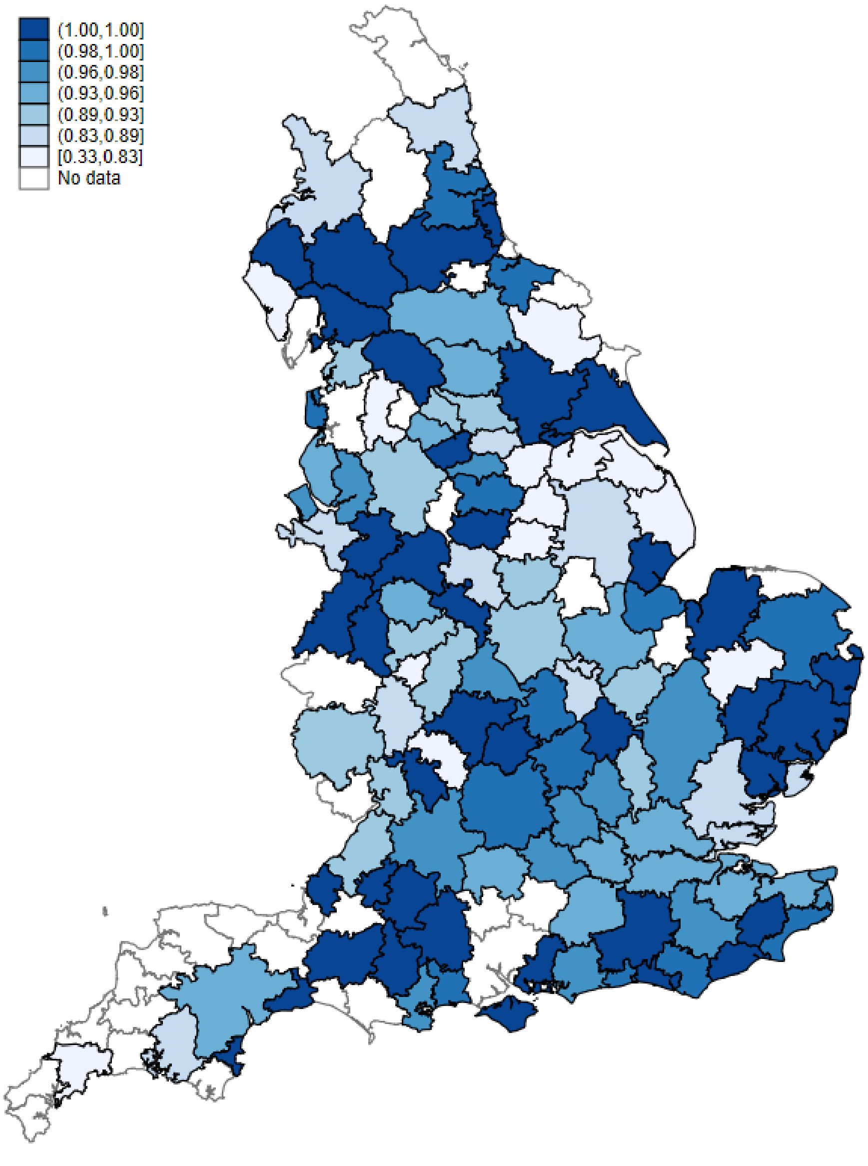

For each TTWA, the map shows the student-weighted share of schools which, in 2009 offered A-level English literature.

Figure 1.A5: Student-weighted share of schools offering A-level history

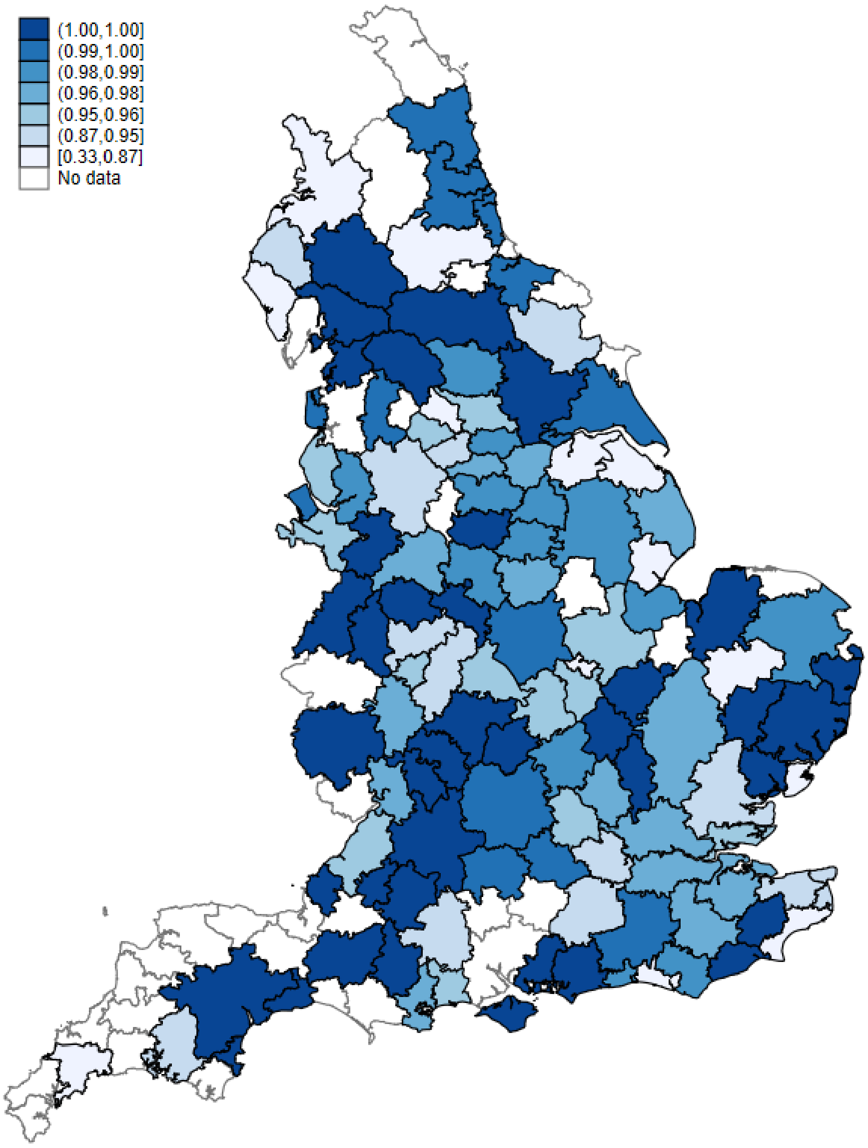

For each TTWA, the map shows the student-weighted share of schools which, in 2009 offered A-level history.

Figure 1.A6: Student-weighted share of schools offering A-level mathematics

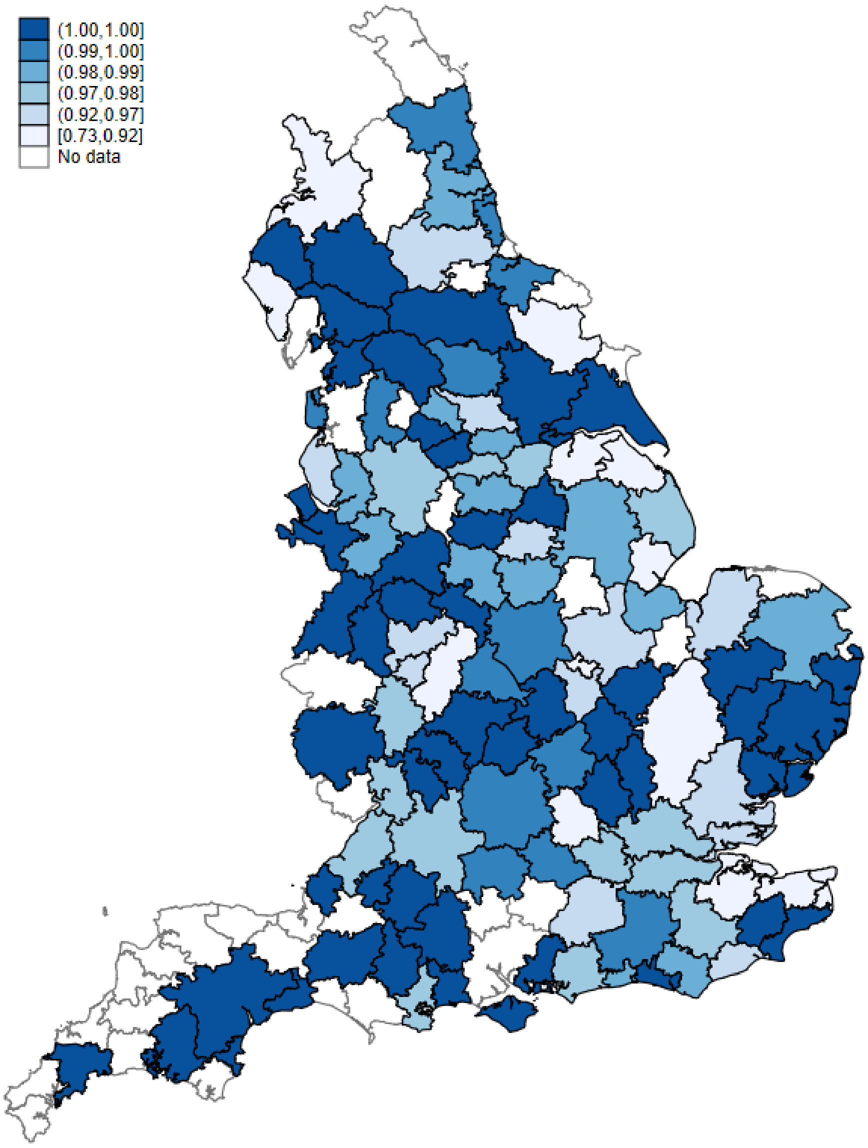

For each TTWA, the map shows the student-weighted share of schools which, in 2009 offered A-level mathematics.

Figure 1.A7: Student-weighted share of schools offering A-level further mathematics

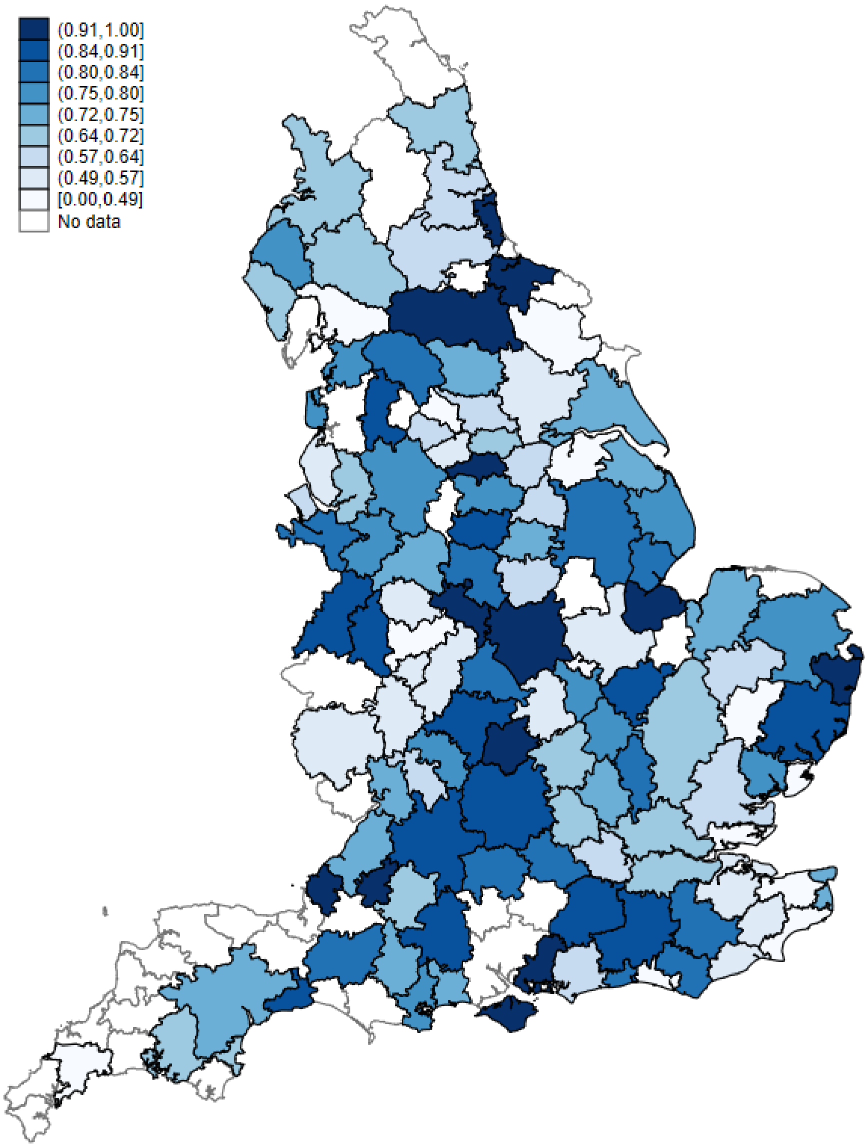

For each TTWA, the map shows the student-weighted share of schools which, in 2009 offered A-level further mathematics.

Figure 1.A8: Student-weighted share of schools offering A-level physics

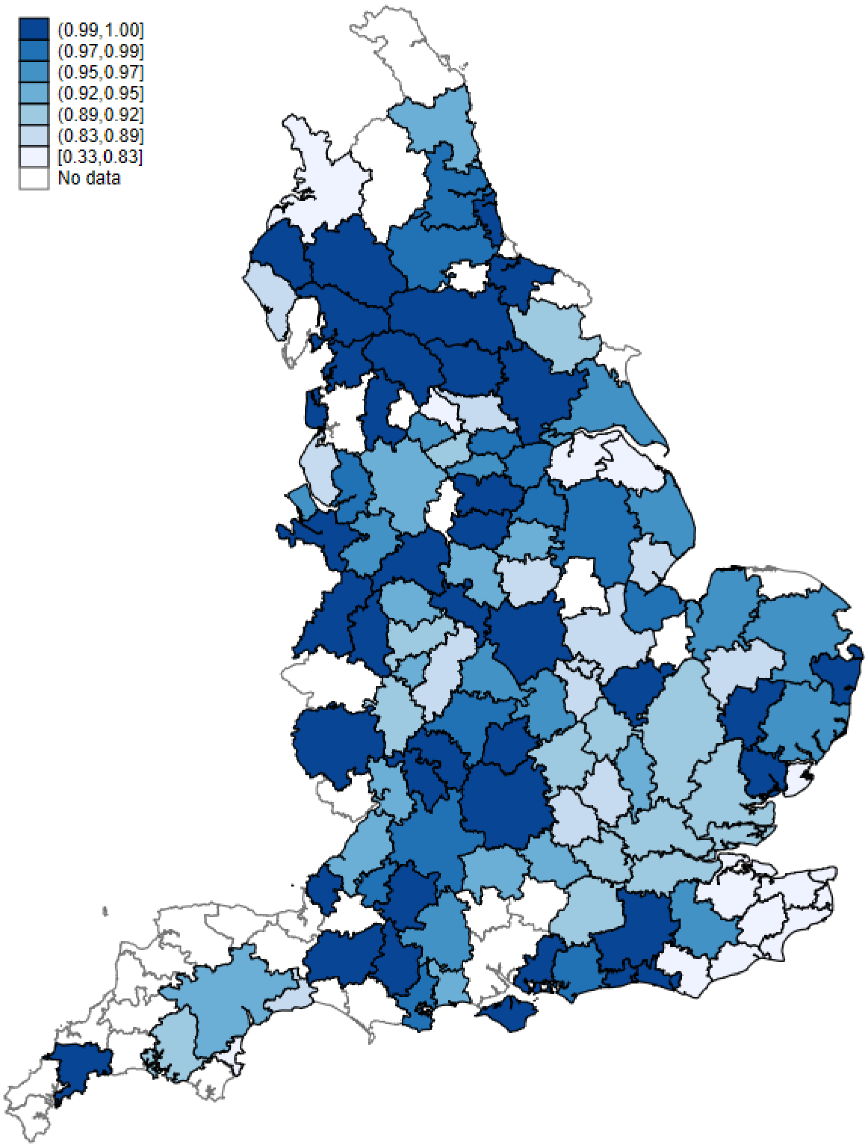

For each TTWA, the map shows the student-weighted share of schools which, in 2009 offered A-level physics.

Table 1.A1: Relationship between local total employment and student degree enrolment

	A-Level Enrol.	A-Level Passes	Uni Enrol.	A-Level Enrol.	A-Level passes	Uni Enrol.
Indexed total emp.	-0.0657**	-0.0738	-0.0590*			
	(0.0247)	(0.0962)	(0.0260)			
Δ_5 Indexed total emp.				-0.0534	-0.0894	-0.0419
				(0.0293)	(0.0838)	(0.0268)
Year FE	Y	Y	Y	Y	Y	Y
N	1705	1705	1705	930	930	930

Data at the TTWA x Year level. Employment shares and changes from the BSD. Student information from the NPD. Cohorts that were in Year 10 from 2004-2013 are included. Student information for HE enrolment comes from HESA. For university enrolment, cohorts that were in Year 10 from 2005-2012 are included. Students are allocated to the first degree of the highest level in which they enrol within 9 years of Year 10.

Appendix 1.B Data Sources

Subsection 1.B.1 Business Structure Database (BSD)

To observe local variation in skill demand, I use the ONS Business Structure Database (BSD), an annual snapshot of the Inter-Departmental Business Register (IDBR), which is a live register of data collected by HM Revenue and Customs via VAT and Pay as You Earn (PAYE) records. This covers all businesses liable for VAT (turnover exceeds the VAT threshold) and/or has at least one member of staff registered for the PAYE tax collection system, then the business will appear on the IDBR (and hence in the BSD). My primary unit in this data is the Local Unit (establishment), which is tied to a specific Enterprise (firm) in each year. I use the location, industry, and employment for each establishment from 2002-2019 to compute the local industrial composition and map that industrial composition to demand for skills using the LFS, as described below.

Subsection 1.B.2 Student Data

Data on students and education choices comes from several datasets provided by the Department for Education (DfE). The main source of data, from which I use information on students' geographic location in the specified base year, is the National Pupil Database (NPD). This is a linked administrative dataset of all students in state-funded schools in England. I observe students' location, demographic characteristics, school information, and exam results during the period of compulsory education. I also use data from the Higher Education Statistics Agency (HESA) dataset, which covers students registered at a higher education provider in a course that leads to a qualification at level 3 or higher.

Subsection 1.B.3 Labour Force Survey (LFS)

I use the UK Labour Force Survey (Office For National Statistics 2023) to map from industries to skills. To define the skill composition, I harmonize the EHATFLD (field) and HIQUAL (level) variables over time, and populate them for earlier years using the given information. The data used is for those ages 18 to 34, employed, in years 2004-2008, quarterly. The field of education is assigned for those who completed a qualification at Level 3 or higher. While A-Levels are a level 3 qualification in the qualifications framework, field is not assigned for students taking A-levels.

I consider all qualifications for which a field is assigned; specifically the field shares include those in both Level 3+ vocational qualifications, which results in the skill mix being somewhat different than those typically taken in academic qualifications. I do this as the fields studied in vocational qualifications are still indicative of the relevant academic qualifications.

The fields and level of qualifications I consider are listed below:

Code	**Label**
0000	Generic programmes and qualifications
0100	Education
0200	Arts and humanities
0300	Social sciences, journalism and information
0400	Business, administration and law
0500	Natural sciences, mathematics and statistics
0600	Information and communication technologies
0700	Engineering, manufacturing and construction
0800	Agriculture, forestry, fisheries and veterinary
0900	Health and welfare
1000	Services

Table 1.B2: Qualification Fields

Code	**Label**
0	Other
1	Degree or equivalent (Level 6 or above)

Table 1.B3: Qualification Levels

Appendix 1.C Data Processing

Subsection 1.C.1 Mapping Fields

After using the LFS to map the fields for workers, I standardize the DfE and HESA data to the same fields. A-Level qualifications are assigned a field based on the sector / subject framework values assigned to the qualification in the Qualification Accreditation framework. These values provide detailed description of the subject studied (e.g. Accounting and Finance; Transportation Operations and Maintenance; Hospitality and Catering), which allows me to map them to the standardized fields used in the LFS. For example, A-levels in economics are mapped to 0300 - Social sciences, journalism and information, and biology is mapped to 0500 - Natural sciences, mathematics and statistics.

University degrees are assigned a subject based on the Joint Academic Coding System (JACS) to match the assignments in the LFS. The exact mapping used in the LFS is spelled out in the data dictionary, and I use the JACS3 codes provided in the HESA data to map degrees (by title) accordingly.

Subsection 1.C.2 BSD Smoothing

Despite being an administrative database, the BSD has some gaps in reporting. To avoid erroneous interpretation of missing data as a decrease and immediate increase in industry employment (and therefore the associated skills), I use the imputation procedure described below. For firms with a gap in reporting between two years in which they are active, I assume their location and industry stays the same as in the initial year. I linearly impute employment between the start and end year, as in the example below.

Year	Employment	Imputed year	Imputed year
2002	210	2002	210
		2003	220
		2004	230
2005	240		

Table 1.C4: Imputation example. The location and industry from the 2002 record would be used for the 2003 and 2004

Chapter 2

Familial persistence of geographic mobility

2.1 Introduction

Geographic mobility represents a critical mechanism for economic opportunity and labour market adjustment in the United States. However, migration rates have declined substantially in recent decades, prompting concerns about diminished economic dynamism and individuals' capacity to access better opportunities across regions. While there has been extensive research on the individual-level determinants of mobility—such as education, housing markets, and local economic conditions—less attention has been paid to how mobility patterns might persist across generations within families.

In this paper, I consider how family history shapes individual geographic mobility decisions along several dimensions. I develop a novel measure of familial mobility: whether an individual was born in the same state as their parents. Using data from the Panel Study of Income Dynamics (PSID) (Survey Research Center 2020), I document significant differences in short-term, medium-term, and lifetime migration rates between adults who were born in the same state as both parents and those who were not. These differences persist after controlling for education, demographics, and time-varying characteristics like employment status and homeownership. And they are of similar magnitudes regardless of gender or education. I further decompose moves into different types—moves away from home states, return migration, and repeat migration— to better understand the patterns that may lead to aggregate differences.

Familial persistence of geographic mobility has both aggregate and individual implications. First, migration is an important labour market adjustment mechanism for both individuals and communities. The main adjustment comes through a decrease in in-migration (Monras 2018). If, as I find, mobility is persistent within families, selection on propensity to move would leave only those (families) with a preference to stay behind in disadvantaged locations. Further, reduced in-migration would not refresh the stock of mobile families in undesirable locations.

While the data used in this project does not allow for detailed location analysis, these patterns have potentially significant implications for the spatial distribution of propensity to move or 'preference for place' as commonly modelled. Consider the spatial sorting model of Diamond and Gaubert (2022), which assumes location-taste preference shocks are identically distributed across space. The results I find can be modelled as intergenerational transmission of preference for place. In a GE setting, over time, the selection mechanism described above would lead these locations to have distribution of location-taste preference shocks different from more desirable locations. This also has implications for place-based redistribution.

This work also has implications for the political divides seen throughout the country and across the world. These divides have recently been characterised as a rift between mobile individuals who do not see location as a barrier from opportunity and those who derive their sense of place from a specific location (Goodhart 2017). In this context, my results can be interpreted as the transmission this characteristic.

This work contributes to several areas involving spatial sorting and geographic mobility. It ties together different threads in the existing literature on family ties and geographic mobility: migration history and family location. Existing work on this topic typically focuses on one or the other, finding that those who live near family are less likely to migrate (Ermisch and Mulder (2019), Mulder and Malmberg (2014), Mulder and Gillespie (2024)). There is also literature documenting the impact of caring responsibilities on migration decisions, which finds increased mobility for those with siblings, compared to only children (Rainer and Siedler 2009). I find that moves of all types are more likely for individuals whose parents moved before they were born.

The existing literature on the impact of early mobility has focused on moves during childhood. Childhood moves have two effects: the child gains 'experience' moving and their existing extra-familial ties are broken. Generally, moving during childhood results lower adult income particularly for students who moved in adolescence (Tønnessen, Telle, and Syse 2016). This association is increasing in the number of moves. In contrast, my research design focuses on individuals whose parents moved before they were born, meaning they were born into a household with experience moving, but do not necessarily have their ties severed or disruption during childhood. The impact of parental migration is documented not only by economists, but also be demographers and sociologists, with studies finding positive relationships between parental mobility and education (Wassink and Viera 2021).

In contrast, recent work on intergenerational mobility emphasises the importance of place in social mobility. Children who moved to lower-poverty areas before age 13 have much better long-run outcomes, in terms of college attendance and lifetimes earnings, compared to those who do not move (Chetty, Hendren, Kline, and Saez 2014). Families who are immobile therefore may not be able to take advantage of opportunities in other areas. Furthermore, this effect can compound over time with 'mobile' families having both a better initial location endowment and higher propensity to move.

There is a large literature documenting on the individual and aggregate factors that contribute to geographic mobility. Causal evidence indicates that education is one of the strongest predictors of interstate mobility, as a result of less location-specific human capital and geographically broader labour market opportunities (Malamud and Wozniak 2012). Individual aggregate housing conditions also play an important role, with homeowners being tied to their current location (Winkler 2011). My research sheds light on an underexplored dimension of heterogeneity in these findings. Controlling for the individual characteristics known to contribute to geographic mobility, I find that parental mobility still plays a significant role, and that the effect is present for all groups.

My findings also have implications for the literature on spatial sorting. The feedback loop created by the segregation of high-skilled, high-productivity workers and firms would be exacerbated by familial persistence in mobility as well as education (Diamond (2016), Gaubert (2018), Diamond and Gaubert (2022)). With persistent migration preferences, mobile, educated, families will concentrate in "superstar cities" with lower-opportunity areas continuing to amass individuals with low motility. This is further compounded by constrained housing supply in high-productivity areas acting as a barrier to migration for lower-skilled workers (Ganong and Shoag 2017).

Finally, this work fits into the broader literature on intergenerational persistence. This phenomenon has been documented for education, earnings, and occupation (Black, Devereux, and Salvanes (2005), Black and Devereux (2011), Ventura (2023)). Much of the recent literature has turned to mechanisms, where persistence has also been found in aspirations and risk preferences (Kimball, Sahm, and Shapiro (2009), Blanden, Doepke, and Stuhler (2023)). Geographic mobility may be influenced by these economically important factors, and also has implications for earnings and other economic outcomes.

The paper proceeds as follows: Section 2.2 describes the data and factors associated with parental mobility. Section 2.3 describes the empirical approach and interpretation. Section 2.4 presents the main results on the relationship between parental mobility and adult children's migration patterns, and heterogeneity by individual characteristics and move type.

2.2 Data and descriptive statistics

2.2.1 PSID Data

My main source of data is the PSID, obtained from for Social Research (2020). The data begins in 1968 with a nationally representative "core" sample of nearly 18,000 individuals from about 5,000 families. In addition to the families included in 1968, a refresher batch of immigrants was added in 1997 to maintain national representativeness.

The study follows these individuals and their descendants, with the most recent data from 2017 including responses from 24,000 individuals form 10,000 family units. Interviews are conducted annually through 1997 and biennially thereafter. Importantly, this dataset tracks families and relationships over time, allowing me to observe not only an individual's move history, but also that of their family.

In each survey year, I observe each respondent's current state, which allows me to construct indicators of whether the individual has moved in the last year, or longer. Additional geographic data, including birth state for all respondents, as well as a rich selection of covariates is available relating to individual observable characteristics as well as family characteristics.

2.2.2 Mobility measures

Much of the existing literature focusing on the US relies on population level cross-sectional data. In these datasets, measures of longer-term mobility are often limited to an indicator of whether an individual lives in the same state where they were born or where their child was born, ignoring any intermediate moves. This understates migration rates, and the panel data set allows me to set a bound on by what extent.

In later years, the panel moves to only biennial sampling, thus some interpolation is necessary. This understates the true migration rate in these years for two reasons. First, it does not account for return migration during the two-year period. In years where annual sampling took place, a move away from a state in one year and a move back in the following year would be observed and counted as a move in each of these years; in years with biennial sampling, the move would not be observed at all. Second, repeat moves in consecutive years are observed only once. This understatement affects both the measures of migration common in the literature, and the more detailed ones I construct which consider intermediate moves.

In addition to migration rates for individuals during the panel, I use historical information to construct mobility rates for parents. From 1970 onward, the birth and childhood states of the parents of household heads is collected; beginning in 1997, this is also collected for wives of household heads. I can thus tell whether heads and their wives, in addition to children in the household, were born in the same state as their parents. For children born into the panel, I have parental demographic information; however, detailed parental data (for example, year of birth) is not available for those whose parents are not in the PSID sample. When considering whether a parent has moved, I do consider moves originating outside of the United States.

My main specifications define parental mobility as an indicator of whether an individual was born in a different state than at least one of their parents. I use this indicator, rather than childhood mobility, to avoid the repeat migration channel, as those who have moved once in their lives are more likely to continue moving (Borjas (2001), Constant and Zimmermann (2012), DaVanzo (1983)). In some specifications, to better relate to the literature, I also include information on childhood moves. This is derived from childhood state in the panel, when available, or from observed states for those who were in the panel as children.

2.2.3 Sample selection and weighting

As the focus of the paper is interstate migration, I restrict my analysis to individuals who live within the US. For each mobility measure, they must also live within the US at the beginning of the mobility period. As I consider only moves within the US when calculating interstate migration rates, an immigrant will not have an indicator of lifetime mobility but will for 1-year and 5-year mobility.

I also restrict my sample to working-age individuals (ages 18-65). This is done for two reasons. First these are the individuals for whom mobility is most likely to influence immediate labour market outcomes. Furthermore, as parental move history is my key variable of interest, including children who live with their parents would include decisions whether to move that were instead made by the parents, confusing interpretation.

2.2.4 Additional Data Sources

To account for heterogeneity in local factors influencing migration rates, I include state-level controls for land area and unemployment. I use state land area from U.S. Census Bureau (2020)[1]. Unemployment rates come from the BLS Local Area Unemployment Statistics, and are available from 1976 onward (U.S. Bureau of Labor Statitics 2020). The rates used are the annual averages calculated by the BLS for the civilian noninstitutional population. As with all other time-varying variables, the rate used is that in the 'origin' state and year (i.e. 1 or 5 years before).

1. The land area is as of 2010, but remains stable over time during the course of this study.

2.2.5 Descriptive differences

Table 2.2 compares working-age individuals with at least one parent who moved before they were born (children of movers) to those born in the same state as both of their parents (children of stayers). Children of movers, shown in column 1, are different in many observable ways; these differences are nearly all in the direction that is associated with higher geographic mobility. On average, they are more educated, which is positively correlated with geographic mobility (Malamud and Wozniak 2012). They are also slightly younger (i.e. born later in the sample). On average they live in slightly larger states. They are much more likely to have moved as children, consistent with repeat migration by their parents.

Moving to time-varying characteristics, those born to mobile parents are less likely to own a house, be employed, or have children. As with the individual characteristics, the direction of these differences is such that those born to mobile parents have characteristics associated with higher geographic mobility. Looking at their recent mobility, they've spent fewer years in their current state, are less likely to have lived there long-term ($\geq$ 10 years), and are more likely to live in a state in which they'd not been previously observed.

In some specifications, I add fixed effects for the individuals' extended family; that is, the family originating from the same 1968 interview. This restricts the sample to only families containing some individuals whose parent(s) moved and some whose did not. Table 2.1 compares individuals within those families. Columns (1) through (4) compare the raw differences between those two groups, consistent with table 2.2. Columns (5) and (6) show their results of a regression of each characteristic on parental mobility and family fixed effects.

Columns (1) and (2) show that the sample means for this restricted sample are relatively similar to those for the full sample shown in table 2.2. Comparing columns (3) and (4) with table 2.2, the differences between those whose parents moved and those who didn't for these families in terms of demographic characteristics broadly follow the same patterns as the main sample. In some cases, the differences are slightly smaller. However, there are key differences in the moving histories of these families, with the number of years in the current state, rate of long term residency, and probability of living in a new state being equal for those whose parents moved and those whose did not. In all cases, these are between the rates in the full sample for those whose parent(s) moved and those whose did not, and closer to the values for those whose parent(s) moved.

Turning to columns (5) and (6), and setting aside recent mobility, adding family fixed effects changes the magnitude of the differences but does not substantively alter the dimensions along which the differences are significant (in either an economic or statistical sense). However, the recent mobility results show a complete reversal from the full sample. Here, compared to other members of their extended family, individuals whose parent(s) moved have lower recent mobility along all dimensions than those whose parents did not.

2.3 Empirical strategy

2.3.1 Regression setup

A key choice in the analysis of geographic mobility is the unit considered. In the context of the impact of mobility on local labour markets, the ideal unit of analysis would be that of inter-labour-market moves. However, even when labour markets can be defined, their evolution over time makes consistent measurement difficult. Furthermore, the data necessary to identify in which labour market an individual lives and has lived is often not accessible.

Due to the availability of variables in the PSID, I focus on interstate migration as my main variable of interest. Compared to local labour markets, the consequences of using this measure are mixed. Much of the existing literature on migration the United States has focused on migration between states and counties. Measuring mobility based on counties overstates labour market mobility, as multiple counties are often part of a single labour market. In contrast, interstate migration can understate geographic mobility though there are factors in both directions. While some labour markets that cross state borders (most prominently, those surrounding New York City and Washington, D.C.), resulting in an overstatement, migration between labour markets within the same state (e.g. between San Francisco and Los Angeles) is not captured by interstate migration measures.

My main specification is as follows:

$$y_{i,t} = \beta M_i + \lambda X_{i,t'} + \gamma u'_{s,t} + \alpha r_s + \delta_t + \varepsilon i,t \tag{2.1}$$

$y_{i,t}$ is the outcome of interest: 1-year, 5-year, or lifetime migration. M_i is a mobility indicator, and the variable of interest. In most specifications, this takes a value of 1 if an individual was born in a different state than their parents. $X_{i,t'}$ is a vector of individual demographic controls in the base (origin) year. It includes highest education and race, which do not change over time, as well as time-varying controls for housing and employment status and quartic in age. $u_{s,t'}$ is origin state unemployment in the base year. Following Rosenbloom and Sundstrom 2004, I include the square-root of origin-state area, r_s (in some specifications, I instead include a state fixed effect). δ_t is a time fixed effect to account for changes in both the data and the migration rate over time.

In other specifications, I also include a family fixed effect. Each individual in the sample is tied to an original (1968) interview family. I refer to individuals associated with the same 1968 interview as being in the same (extended) family. In these specifications, I use variation within families; thus, the sample is restricted to only (extended) families with variation in parental mobility. Differences within families can come both across and between generations[2].

2. To fix ideas, a few examples of within-family variation: cousins where some have a parent who moved and other don't (either by marriage or different mobility patterns from their parents who are siblings); siblings where one was born before their parents moved (or returned) and one was born after; an adult born to immobile parents who moves then has a child.

2.4 Results

First, I consider all moves, comparing those who were born in the same state as both their parents to those who were not. Figure 2.1 shows the evolution of short- and long-term geographic mobility over the last half-century, split by parental mobility. For every duration of move, the migration rates for those born to mobile parents are higher than those born in the same state as their parent(s). The trends for both groups over time are also the same, with increasing lifetime mobility but decreasing short-term mobility over time. I discuss the types of moves and patterns there in the subsequent sections.

As seen in section 2.2, those whose parent(s) moved are different along many dimensions than those whose parent(s) did not. I control for these factors in the regressions shown in table 2.3, which shows the differences in 1-year, 5-year, and lifetime migration rates between those whose parent(s) moved and those whose parent(s) did not. The first column shows the raw difference in migration rates, with those born to mobile parent(s) about 30 percent more likely to move across all durations.

In the next column, I control for year fixed effects. In addition to differences in migration rates over time, this also helps address concerns about changes in the nature of available data as the panel progresses. Early in the study period, many individuals are missing parental information. Later generations included in the PSID have both a richer individual migration history and more family migration information available. Despite these differences, adding these fixed effects does not meaningfully alter the results.

I then add controls to address the differences seen in table 2.2. I first add controls for education and demographics, shown in column 3. I then add time-varying characteristics: a quartic in age, employment status, and homeownership, shown in column 4.

While controlling for these characteristics reduces the magnitude of the effect, the results remain both statistically and economically significant. With controls, the short-run migration rates are 20-25 percent higher for those whose parent(s) moved, and the lifetime mobility is 50 percent higher. In column 5, I include controls for the land area and unemployment rate in the origin state. These do not meaningfully alter the results, and this is the preferred specification I refer to later. I also confirm in column 6 that using a state fixed effect, rather than land area, does not meaningfully alter the results.

Finally, I add a fixed effect for the extended family[3] to which the individual belongs. Thus, column 7 compares individuals within the same 'family tree', distinguishing those from mobile versus immobile branches. Here, the magnitude of the effects for the short- and medium-moves shrinks to be indistinguishable from 0. This implies that the differences in short- to medium-run migration rates seen in column 6 come from differences between 'mobile' and 'immobile' families. This is consistent with the offsetting demographic and move history differences seen in table 2.1. However, individuals whose parent(s) moved are still significantly more likely to live in state different than the one in which they were born. This could be due to a higher chance of migration during childhood (panels A and B consider only working-age adults), or a decrease in repeat or return migration.

3. The initial interview family to which the individual is associated.

2.4.1 Changes over time

Much of the recent literature on migration in the US focuses on the declining rates of geographic mobility since the mid-20th century. In this section, I examine whether the effect I see changes over time. Table 2.4 interacts the parental mobility indicator with an indicator for each decade of the PSID. For short and medium-run mobility, the effect is indeed larger in the first decade of the PSID, but it remains both economically and statistically significant in later years. The magnitude of the impact on 1-year move rate declines somewhat in the later part of the study, while the impact on 5-year moves is relatively stable for all decades after the 1970s. Thus, changes in familial persistence do not help to explain the overall decline in mobility[4].

2.4.2 Heterogeneity

I now consider how the results differ by group. I consider heterogeneity along two dimensions: sex and education. In making these comparisons, I consider the preferred specification, with full controls, with and without family fixed effects.

I first consider differences between those who have only a high school education or less and those who continued their studies after high school. The evolution of 5-year move rates for these groups over time is shown in figure 2.2. In all years, those who studied beyond high school are more mobile than those who did not, with even the more-educated children of stayers moving more often than less-educated children of movers. This points to the importance of other factors identified in the literature in the moving decision. However, within-group there remain significant differences based on parental mobility.

The first two columns of table 2.5 show the overall results (columns 5 and 6 of table 2.3 for reference). Columns 3 and 4 show the same regressions but with the sample restricted to those have only a high-school education, and columns 5 and 6 show are those who studied beyond high school.

Comparing columns 3 and 5, those with college are much more likely to move across all durations. Although the effect size for the lower education group, shown in column 3, is smaller, the relative magnitude is similar. As in the aggregate, the 1-year migration rates increase by 20-25 percent for both groups. Over 5 years the effect is 30 percent for those with high school compared to only 20 percent for those with more education.

The results with family fixed effects, shown in columns 4 and 6, compared individuals to others in their (extended) family with the same education level. In both cases, these show the same pattern as the aggregate results: within families, those born to mobile parents are not significantly more likely to move than those born in the same state as their parents.

Columns 7 and 9 show the results with the sample restricted to men and women, respectively. The short-run mobility impact for women is slightly larger, both in terms of absolute and relative magnitude, though the lifetime coefficient is somewhat smaller. This discrepancy can be explained by fewer repeat moves, or higher rates of return migration. Turning to column 8, which compares men to the other men

4. The decade fixed effects, not shown in the table, are consistent with declining short and medium-run mobility.

in their family, the results again show the same pattern as the full-sample results.

Column 10 compares women whose parent(s) moved to other women in their family whose parent(s) did not. Unlike the other within-family results, the magnitude of the difference in short-run mobility for this group is of the same order of magnitude as the results without family controls. For both 1-year and 5-year moving rates, women whose parent(s) moved are 10-15 percent more likely to move than those whose parent(s) did not. The differences here could be driven by the importance of caring responsibilities (Compton and Pollak 2014).

2.4.3 Move types

I now examine different types of moves. Figure 2.3 decomposes 5-year moving rates into moves away from home, moves back home, moves while away, and moves between home states. Here, home states are the places where the individual was born or lived as a child.

Panel A shows this decomposition for adults who were born in the same state as their parents. The rate of moving away from their home state declined from 1970s to the 1990s and has remained relatively stable around 3-4 percent since then. Panel B shows the same composition for those born in a different state than at least one of their parents. The rate of moving away is consistently higher than for those born in the same state as their parents, but has also slowed, comparing the years before and after the mid-nineties. For both groups, the repeat migration rate (moves when already away) has remained stable over time, with those born to mobile parents moving about 50 percent more. In absolute terms, return migration rates are also higher for those born to mobile parents, however a larger share of these individuals live away from their home state(s), so the relative magnitudes are similar.

An important aspect of geographic mobility, when considering economic dynamism and long-run differences in local characteristics, is the propensity of individuals to leave the conditions which they have inherited. Table 2.6 restricts the sample to individuals who have never left their birth state. Over all durations, the average mobility is about 40 percent lower for this group than the full sample. I chose to restrict to this group, rather than to those who have never moved in their lifetime, to include those who were moved by their parents as children. This allows the results to be interpreted as the impact of the parental migration decisions (for their young family) on the individuals' later choices (moves as adults).

As in the full sample, with full controls, the migration rates for those born in a different state than their parent(s) are about 25 percent higher than the rate for those born in the same state as both parents. The final column of the table compares individuals to others in their (extended) family who have not moved as adults. As in the full sample, there is no significant difference.

2.4.4 Move history

There is a significant literature in the individual propensity to migrate, focusing on individual rather than familial moving history. To better relate my findings to this literature, I consider and compare the impact of childhood moves to family move history. Table 2.7 adds an indicator for whether the individual moved before age 18, as well as an interaction term, to the main specification shown in table

2.3.

Parental moves before a child was born remain important for their future mobility, and the effect is additive with additional moves during childhood. The impact of parental moves pre-birth is roughly half the impact of the impact of moving as a child. Furthermore, the magnitude of the effect is about half of the magnitude without childhood move controls – it remains both economically and statistically significant.

These specifications help to elucidate some of the possible mechanisms behind the increased geographic mobility. Children of parents who moved states before they were born but not during childhood would likely have a stable social peer group[5], but family ties in other locations.

2.5 Conclusion

Adults born in a different state than their parents were are much more likely to move throughout their lifetime. While these individuals have different observable characteristics than those who were born in the same state as their parents, the differences in migration rates are robust to controls for these factors. Furthermore, these differences - about a 20 percent increase in migration rates, are consistent across subgroups by education and gender. The effects also consistently present throughout the period I study.

Within (extended) families, there are not significant differences in migration rates between those born in different states than their parents and their cousins and other family members born in the same state as their parents. This points toward extended family ties as an important factor in migration decisions. The exception to this is that women born in a different state than their parent(s) remain more likely to move, even compared to their extended family.

When examining different types of moves, those born to mobile parents exhibit higher rates of all move types: moves away from home states, return migration, and repeat migration. Even among individuals who have never left their birth state as adults, those with mobile parents have approximately 25 percent higher migration rates, suggesting a persistent difference in propensity to relocate.

These results have implications for the distribution of location preference (or moving cost) in spatial sorting models, as those with a lower preference for their initial location move away, and subsequent generations' preferences are related to those of the current generation. An interesting modification to existing models would be an extension to include this property of preferences.

While this research establishes the importance of familial mobility history, several questions remain for future research. First, further exploration of the mechanisms driving this persistence, such as family ties, risk preferences, information dissemination or other factors, is relevant to the implications. In particular, the relationship between these factors may be important for diverging economies within countries. Furthermore, additional work can be done to understand the implications of the difference in geographic mobility for other outcomes. For example, how these children differ in their education choices, or their response to economic shocks. While this work identifies a clear mobility pattern, there is much more to be done to fully understand the patterns described here.

5. Here I consider only inter-state moves, but it is possible that the families are more mobile intra-state, lessening social ties.

Bibliography for Chapter 2

Black, Sandra E, and Paul J Devereux. 2011. "Recent developments in intergenerational mobility." *Handbook of labor economics* 4:1487–1541.

Black, Sandra E, Paul J Devereux, and Kjell G Salvanes. 2005. "Why the apple doesn't fall far: Understanding intergenerational transmission of human capital." *American economic review* 95 (1): 437–449.

Blanden, Jo, Matthias Doepke, and Jan Stuhler. 2023. "Educational inequality." In *Handbook of the Economics of Education,* 6:405–497. Elsevier.

Borjas, G.J. 2001. "Migration, Economics of." In *International Encyclopedia of the Social Behavioral Sciences,* edited by Neil J. Smelser and Paul B. Baltes, 9803–9809. Oxford: Pergamon.

Chetty, Raj, Nathaniel Hendren, Patrick Kline, and Emmanuel Saez. 2014. "Where is the land of opportunity? The geography of intergenerational mobility in the United States." *The quarterly journal of economics* 129 (4): 1553–1623.

Compton, Janice, and Robert A Pollak. 2014. "Family proximity, childcare, and women's labor force attachment." *Journal of Urban Economics* 79:72–90.

Constant, Amelie F, and Klaus F Zimmermann. 2012. "The dynamics of repeat migration: A Markov chain analysis." *International Migration Review* 46 (2): 362–388.

DaVanzo, Julie. 1983. "Repeat migration in the United States: who moves back and who moves on?" *The Review of Economics and Statistics,* 552–559.

Diamond, Rebecca. 2016. "The determinants and welfare implications of US workers' diverging location choices by skill: 1980–2000." *American Economic Review* 106 (3): 479–524.

Diamond, Rebecca, and Cecile Gaubert. 2022. "Spatial sorting and inequality." *Annual Review of Economics* 14:795–819.

Ermisch, John, and Clara H Mulder. 2019. "Migration versus immobility, and ties to parents." *European Journal of Population* 35 (3): 587–608.

Ganong, Peter, and Daniel Shoag. 2017. "Why has regional income convergence in the US declined?" *Journal of Urban Economics* 102:76–90.

Gaubert, Cecile. 2018. "Firm sorting and agglomeration." *American Economic Review* 108 (11): 3117–3153.

Goodhart, David. 2017. *The road to somewhere: The populist revolt and the future of politics.* Oxford University Press.

Kimball, Miles S, Claudia R Sahm, and Matthew D Shapiro. 2009. "Risk preferences in the PSID: Individual imputations and family covariation." *American Economic Review* 99 (2): 363–368.

Malamud, Ofer, and Abigail Wozniak. 2012. "The impact of college on migration evidence from the Vietnam generation." *Journal of Human resources* 47 (4): 913–950.

Monras, Joan. 2018. "Economic Shocks and Internal Migration." *Quarterly Journal of Economics* 133 (3): 1209–1253.

Mulder, Clara H, and Gunnar Malmberg. 2014. "Local Ties and Family Migration." *Environment and Planning A: Economy and Space* 46 (9): 2195–2211.

Mulder, Clara H., and Brian Joseph Gillespie. 2024. "Moving and staying in the context of the family: A review and an introduction to the Special Issue." *Population, Space and Place* 30 (3): e2712.

Rainer, Helmut, and Thomas Siedler. 2009. "O brother, where art thou? The effects of having a sibling on geographic mobility and labour market outcomes." *Economica* 76 (303): 528–556.

Rosenbloom, Joshua L, and William A Sundstrom. 2004. "The decline and rise of interstate migration in the United States: Evidence from the IPUMS, 1850–1990." In *Research in economic history,* 289–325. Emerald Group Publishing Limited.

Survey Research Center. 2020. *Panel Study of Income Dynamics, public use dataset.* Produced and distributed by the Survey Research Center, Institute for Social Research, University of Michigan. Ann Arbor, MI.

Tønnessen, Marianne, Kjetil Telle, and Astri Syse. 2016. "Childhood residential mobility and long-term outcomes." *Acta Sociologica* 59 (2): 113–129.

U.S. Bureau of Labor Statitics. 2020. *Local Area Unemployment Statistics.*

U.S. Census Bureau. 2020. *State Area Measurements and Internal Point Coordinates.*

Ventura, Maria. 2023. *Following in the Family Footsteps: Incidence and Returns of Occupational Persistence.*

Wassink, Joshua Thomas, and Janelle Ashley Viera. 2021. "Does parental migration during childhood affect children's lifetime educational attainment? Evidence from Mexico." *Demography* 58 (5): 1765–1792.

Winkler, Hernan. 2011. "The effect of homeownership on geographic mobility and labor market outcomes." *Available at SSRN 1724455.*

Appendix

Figure 2.1: 5-year move rate decomposition by parental mobilty

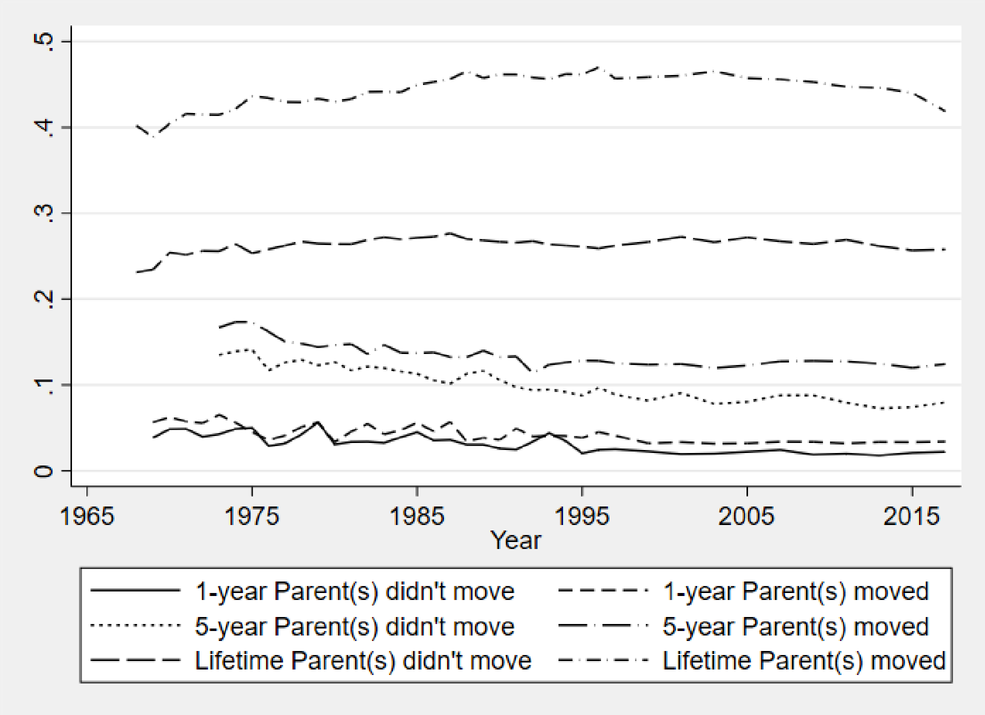

The sample includes all working-age individuals with parental data and is weighted according to individual weights.

Figure 2.2: 5-year move rates by education and parental mobilty

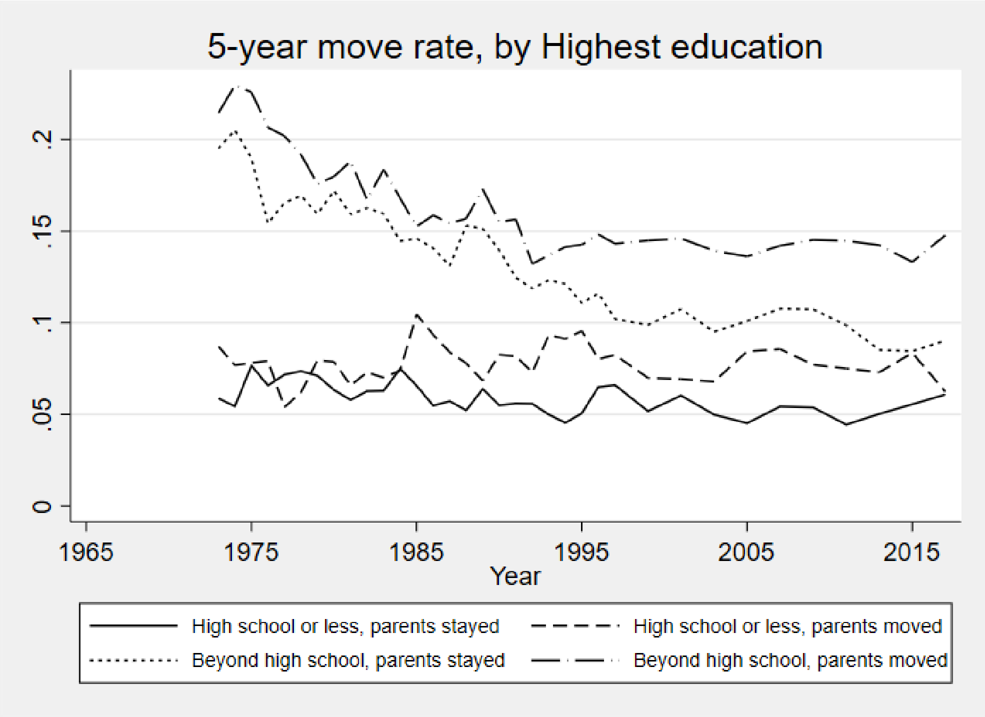

The sample includes all working-age individuals with parental data and is weighted according to individual weights.

Figure 2.3: 5-year move rates by type and parental mobility

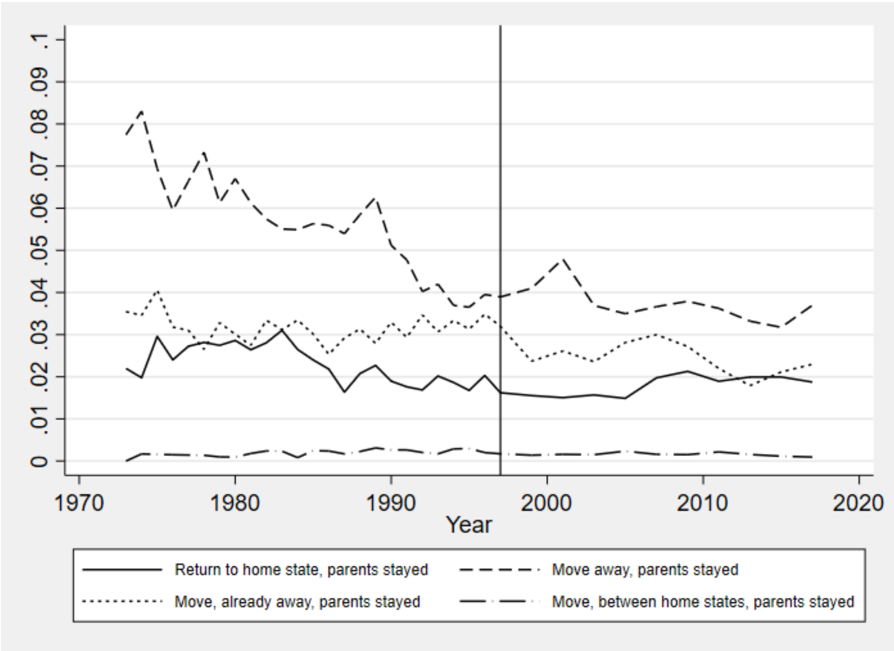

(a) Parents did not move

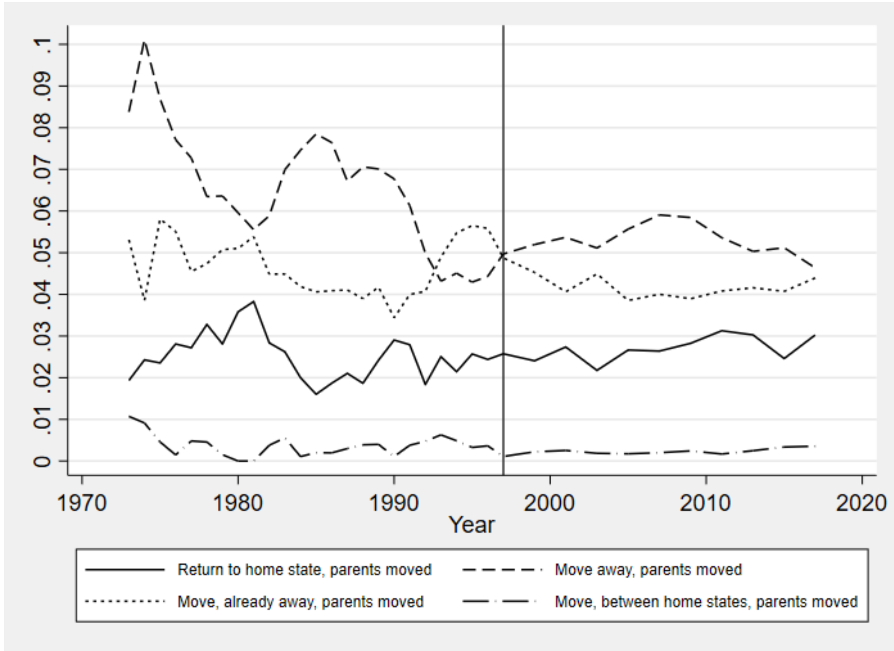

(b) Parent(s) moved

The sample includes all working-age individuals with parental data and is weighted according to individual weights.Panel (a) includes only those who were born in the same state as their parent(s); panel (b) includes only those born in a different state than their parent(s).

Table 2.1: Balance of characteristics

	Parent(s) moved	**Parent(s) didn't move**		
	Mean	Mean	Difference	p-value
Variables	(1)	(2)	(3)=(1)-(2)	
Individual characteristics				
Highest education				
Less than high school	0.062	0.097	-0.036***	(0.000)
High school only	0.256	0.301	-0.045***	(0.000)
Some college	0.301	0.268	0.033***	(0.000)
College or more	0.382	0.334	0.048***	(0.000)
Demographics				
Female	0.498	0.526	-0.028**	(0.002)
Birth Year	1970	1966	3.245***	(0.000)
White	0.799	0.781	0.018*	(0.018)
Black	0.130	0.151	-0.021**	(0.001)
Other race	0.071	0.068	0.003	(0.508)
Birth state				
Birth state unemp.	0.066	0.067	-0.001*	(0.018)
Birth state size	409.649	367.128	42.521***	(0.000)
Early mobility				
Moved as Child	0.337	0.126	0.211***	(0.000)
Number of individuals	5,177	7,103		
Time-varying characteristics				
Housing status				
Own	0.651	0.688	-0.037***	(0.000)
Employment status				
Employed	0.647	0.664	-0.017***	(0.000)
Unemployed	0.045	0.045	0.001	(0.452)
Children				
Has child(ren)	0.371	0.416	-0.045***	(0.000)
Num. children	0.835	0.917	-0.082***	(0.000)
Recent mobility				
Num. years in state	24.670	27.185	-2.515***	(0.000)
Live $\geq$ 10 years in state	0.788	0.829	-0.042***	(0.000)
Live in new state	0.020	0.014	0.006***	(0.000)
Number of person-years	74,905	105,494		

Notes: The sample includes all working-age individuals with parental data. Results are weighted by the number of students in Year 10 in the TTWA in 2002. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$

Table 2.2: Balance of characteristics - Family sample

	Parent(s) moved	**Parent(s) didn't move**			**within fam**	**within fam**
	Mean	Mean	Difference	p-value	Difference	p-value
Variables	(1)	(2)	(3)=(1)-(2)	(4)	(5)	(6)
Individual characteristics						
Highest education						
Less than high school	0.069	0.099	-0.030***	(0.000)	-0.026***	(0.000)
High school only	0.272	0.303	-0.031***	(0.001)	0.007	(0.455)
Some college	0.304	0.274	0.030**	(0.001)	0.035***	(0.000)
College or more	0.355	0.323	0.031**	(0.001)	-0.017	(0.066)
Demographics						
Female	0.503	0.522	-0.019	(0.059)	-0.012	(0.303)
Birth Year	1971	1967	4.349***	(0.000)	5.596***	(0.000)
White	0.786	0.776	0.010	(0.240)	-0.008	(0.056)
Black	0.140	0.163	-0.023**	(0.002)	-0.002	(0.383)
Other race	0.074	0.061	0.013*	(0.013)	0.011*	(0.011)
Birth state						
Birth state unemp.	0.066	0.067	-0.001*	(0.025)	-0.001	(0.415)
Birth state size	411.959	363.340	48.619***	(0.000)	20.767***	(0.000)
Childhood mobility						
Moved as Child	0.323	0.150	0.173***	(0.000)	0.114***	(0.000)
Number of individuals	4,275	5,451				
Time-varying characteristics						
Housing status						
Own	0.644	0.680	-0.036***	(0.000)	-0.028***	(0.000)
Employment status						
Employed	0.641	0.661	-0.020***	(0.000)	-0.020***	(0.000)
Unemployed	0.049	0.047	0.001	(0.238)	0.002	(0.170)
Children						
Has child(ren)	0.394	0.432	-0.038***	(0.000)	-0.045***	(0.000)
Num. children	0.913	0.971	-0.057***	(0.000)	-0.072***	(0.000)
Recent mobility						
Num. years in state	25.736	25.932	-0.196*	(0.021)	2.389***	(0.000)
Long term res	0.814	0.814	0.000	(0.962)	0.039***	(0.000)
Live in new state	0.016	0.016	-0.000	(0.808)	-0.005***	(0.000)
Number of person-years	58,704	80,056				

Notes: The sample includes all working-age individuals with parental dataResults are weighted by the number of students in Year 10 in the TTWA in 2002. *** $p < 0.01$; ** $p < 0.05$; * $p < 0.10$

Table 2.3: Relationship between parental moves and migration rates

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
A. 1-year							
Parent(s) moved	0.0118	0.0116	0.00922	0.00742	0.00746	0.00725	0.00102
	(0.00157)	(0.00159)	(0.00151)	(0.00144)	(0.00144)	(0.00143)	(0.00164)
Observations	167,603	167,603	167,496	167,493	152,654	152,654	152,654
Mean	0.0326	0.0326	0.0326	0.0326	0.0314	0.0314	0.0314
B. 5-year							
Parent(s) moved	0.0369	0.0364	0.0287	0.0241	0.0252	0.0254	0.00266
	(0.00535)	(0.00538)	(0.00508)	(0.00493)	(0.00504)	(0.00504)	(0.00600)
Observations	134,339	134,339	134,262	132,974	118,452	118,452	118,452
Mean	0.109	0.109	0.109	0.110	0.106	0.106	0.106
C. Lifetime							
Parent(s) moved	0.184	0.184	0.177	0.177	0.190	0.169	0.0883
	(0.0155)	(0.0155)	(0.0152)	(0.0152)	(0.0188)	(0.0189)	(0.0229)
Observations	169,945	169,945	169,900	169,900	36,973	36,973	36,973
Mean	0.349	0.349	0.348	0.348	0.291	0.291	0.291
Year FE		✓	✓	✓	✓	✓	✓
Individual char.			✓	✓	✓	✓	✓
Indiv. variable char.				✓	✓	✓	✓
State controls					✓	✓	✓
State FE						✓	✓
Family FE							✓

Notes: The sample includes all working-age individuals with parental data and is weighted according to individual weights. Standard errors are clustered at the 1968 family level. Individual characteristics are highest education, race, and age. State controls are the unemployment rate in the state at the beginning of the period and the area of the origin state (excluded when state FE are included).

Table 2.4: Parental moves and migration rates over time

	(1)	(2)	(3)	(4)	(5)	(6)
A. 1-year						
1968-1977 × Parent(s) moved	0.00997	0.00991	0.0118	0.0110	0.0113	0.00840
	(0.00504)	(0.00505)	(0.00486)	(0.00472)	(0.00463)	(0.00471)
1978-1987 × Parent(s) moved	0.0108	0.0108	0.0104	0.00798	0.00886	0.00496
	(0.00415)	(0.00415)	(0.00404)	(0.00392)	(0.00387)	(0.00380)
1988-1997 × Parent(s) moved	0.0114	0.0115	0.0105	0.00816	0.00819	0.00232
	(0.00302)	(0.00302)	(0.00291)	(0.00283)	(0.00278)	(0.00286)
1998-2007 × Parent(s) moved	0.0108	0.0108	0.00716	0.00568	0.00526	-0.00317
	(0.00207)	(0.00207)	(0.00200)	(0.00194)	(0.00192)	(0.00224)
2008-2017 × Parent(s) moved	0.0133	0.0133	0.00810	0.00662	0.00574	-0.00402
	(0.00174)	(0.00174)	(0.00166)	(0.00163)	(0.00166)	(0.00197)
Observations	167,603	167,603	167,496	167,493	167,493	167,493
Mean	0.0326	0.0326	0.0326	0.0326	0.0326	0.0326
B. 5-year						
1968-1977 × Parent(s) moved	0.0334	0.0334	0.0396	0.0392	0.0410	0.0285
	(0.0166)	(0.0166)	(0.0159)	(0.0157)	(0.0155)	(0.0156)
1978-1987 × Parent(s) moved	0.0247	0.0246	0.0242	0.0199	0.0230	0.00568
	(0.0114)	(0.0114)	(0.0111)	(0.0108)	(0.0107)	(0.0106)
1988-1997 × Parent(s) moved	0.0305	0.0304	0.0266	0.0206	0.0223	0.000986
	(0.00931)	(0.00931)	(0.00895)	(0.00871)	(0.00865)	(0.00895)
1998-2007 × Parent(s) moved	0.0400	0.0399	0.0302	0.0253	0.0262	-0.00347
	(0.00782)	(0.00782)	(0.00745)	(0.00734)	(0.00726)	(0.00805)
2008-2017 × Parent(s) moved	0.0461	0.0462	0.0299	0.0260	0.0248	-0.0115
	(0.00633)	(0.00634)	(0.00600)	(0.00592)	(0.00597)	(0.00703)
Observations	134,339	134,339	134,262	132,974	132,971	132,971
Mean	0.109	0.109	0.109	0.110	0.110	0.110
Year FE		✓	✓	✓	✓	✓
Individual char.			✓	✓	✓	✓
Indiv. variable char.				✓	✓	✓
State FE					✓	✓
Family FE						✓

Notes: The sample includes all working-age individuals with parental data and is weighted according to individual weights. Standard errors are clustered at the 1968 family level. Individual characteristics are highest education, race, and age. State controls are the unemployment rate in the state at the beginning of the period and the area of the origin state (excluded when state FE are included).

Table 2.5: Heterogenity of impact of parental moves

	All		High School or Less		More than High School		Men		Women	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
A. 1-year										
Parent(s) moved	0.00746	0.00121	0.00469	0.000926	0.00891	0.00254	0.00570	-0.00105	0.00880	0.00344
	(0.00144)	(0.00162)	(0.00208)	(0.00264)	(0.00184)	(0.00228)	(0.00211)	(0.00276)	(0.00195)	(0.00264)
Observations	152,654	152,654	58,602	58,602	94,052	94,052	67,358	67,358	85,296	85,296
Mean	0.0314	0.0314	0.0203	0.0203	0.0371	0.0371	0.0327	0.0327	0.0303	0.0303
B. 5-year										
Parent(s) moved	0.0252	0.00269	0.0193	0.00649	0.0281	0.00194	0.0212	-0.00367	0.0282	0.0156
	(0.00504)	(0.00598)	(0.00725)	(0.00881)	(0.00650)	(0.00870)	(0.00742)	(0.0106)	(0.00671)	(0.00984)
Observations	118,452	118,452	45,014	45,014	73,438	73,438	51,940	51,940	66,512	66,512
Mean	0.106	0.106	0.0635	0.0635	0.128	0.128	0.110	0.110	0.103	0.103
Year FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Ind & state controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Family FE		✓		✓		✓		✓		✓

Notes: The sample in columns (1) and (2) includes all working-age individuals with parental data and is weighted according to individual weights.Columns (3) and (4) restrict the sample to only those with high school or less. Columns (5) and (6) restrict the sample to only those with more than high school. Columns (7) and (8) restrict the sample to men only. Columns (9) and (10) restrict the sample to women only. Standard errors are clustered at the 1968 family level. Individual characteristics are highest education, race, and age. State controls are the unemployment rate in the state at the beginning of the period and the area of the origin state (excluded when state FE are included).

Table 2.6: Impact of parental moves on stayers

	(1)	(2)	(3)	(4)	(5)	(6)
A. 1-year						
Parent(s) moved	0.00693	0.00682	0.00554	0.00478	0.00544	-0.000733
	(0.00113)	(0.00113)	(0.00102)	(0.000996)	(0.00103)	(0.00163)
Observations	109,578	109,578	109,543	109,542	99,013	99,013
Mean	0.0184	0.0184	0.0185	0.0185	0.0176	0.0176
B. 5-year						
Parent(s) moved	0.0272	0.0267	0.0212	0.0186	0.0229	-0.00279
	(0.00477)	(0.00473)	(0.00429)	(0.00422)	(0.00436)	(0.00694)
Observations	87,864	87,864	87,841	87,159	76,947	76,947
Mean	0.0724	0.0724	0.0725	0.0724	0.0698	0.0698
Year FE		✓	✓	✓	✓	✓
Individual char.			✓	✓	✓	✓
Indiv. variable char.				✓	✓	✓
State controls					✓	✓
Family FE						✓

Notes: The sample includes all working-age individuals with parental datawho have never moved as adults and is weighted according to individual weights.Standard errors are clustered at the 1968 family level. Individual characteristics are highest education, race, and age. State controls are the unemployment rate in the state at the beginning of the period and the area of the origin state.

Table 2.7: Impact of childhood moves

	(1)	(2)	(3)	(4)	(5)	(6)
A. 1-year						
Moved as child	0.0155	0.0157	0.0129	0.0118	0.0131	0.00420
	(0.00341)	(0.00345)	(0.00327)	(0.00308)	(0.00308)	(0.00406)
Parent(s) moved	0.00902	0.00871	0.00671	0.00490	0.00514	0.000622
	(0.00170)	(0.00173)	(0.00166)	(0.00159)	(0.00159)	(0.00199)
Parents moved x Child move	0.000389	0.000529	0.000979	0.00149	-0.000227	-0.000232
	(0.00447)	(0.00452)	(0.00427)	(0.00407)	(0.00407)	(0.00505)
Observations	161,859	161,859	161,787	161,784	147,085	147,085
Mean	0.0329	0.0329	0.0329	0.0329	0.0318	0.0318
B. 5-year						
Moved as child	0.0509	0.0515	0.0409	0.0388	0.0451	0.0145
	(0.0112)	(0.0113)	(0.0106)	(0.0104)	(0.0109)	(0.0134)
Parent(s) moved	0.0289	0.0283	0.0219	0.0174	0.0196	0.00199
	(0.00580)	(0.00584)	(0.00557)	(0.00542)	(0.00557)	(0.00697)
Parents moved x Child move	-0.00380	-0.00363	-0.00214	-0.00118	-0.00943	-0.00635
	(0.0151)	(0.0152)	(0.0142)	(0.0139)	(0.0142)	(0.0169)
Observations	131,337	131,337	131,283	130,031	115,647	115,647
Mean	0.110	0.110	0.110	0.110	0.107	0.107
Year FE		✓	✓	✓	✓	✓
Individual char.			✓	✓	✓	✓
Indiv. variable char.				✓	✓	✓
State controls					✓	✓
Family FE						✓

Notes: The sample includes all working-age individuals with parental dataand information about childhood location and is weighted according to individual weights.Standard errors are clustered at the 1968 family level. Individual characteristics are highest education, race, and age. State controls are the unemployment rate in the state at the beginning of the period and the area of the origin state.

Chapter 3

Individual consequences of occupational decline

This chapter is jointly co-authored with Per-Anders Edin, Georg Graetz, Sofia Hernnäs, and Guy Michaels.

3.1 Introduction

What are the career earnings losses that workers suffer when demand for their occupations declines? This question is important for policy debates on responses to technologies that replace workers (Acemoglu and Restrepo 2019), and is relevant for broader discussions on labour market transformations due to technological change (see for instance Brynjolfsson and McAfee 2014, Autor 2015, and Caselli and Manning 2019). New labour-replacing technologies no longer threaten only machine operatives and clerical workers. Self-driving vehicles may reduce the employment of drivers (Campbell 2018), and artificial intelligence software challenges professionals such as lawyers and financial investors (Susskind and Susskind 2015) and even fashion designers (Scheiber 2018). This is causing considerable angst. It is therefore important to understand how costly occupation-replacing technologies are for workers, since this informs our thinking about individual welfare, inequality, and human capital investments. It is also important for public policy decisions on taxation, redistribution, retirement, and education, and may even have broader political consequences (Marx 1867; Caprettini and Voth 2017; Bo et al. 2022).

In this paper, we investigate the consequences for the career earnings and employment of individual Swedish workers of large declines in demand for their occupations, which are driven by technological change. We assemble high-quality population-level administrative data spanning several decades with a highly detailed occupational classification and a rich set of control variables. Using these, we regress workers' career earnings on an indicator for (large) occupational declines, controlling for potential confounders as discussed below. We show that the coefficient on the occupational decline indicator can be decomposed into the mean effect of occupational decline on workers whose occupation declines relative to those whose occupation does not, plus a selection term.

This selection term represents the mean difference in outcomes between workers whose occupation declines and other workers in absence of any occupational decline. To address this selection problem, we show that adding individual-level controls addresses individual-level selection on cognitive and non-cognitive skills and parental backgrounds. We also control for occupation-level characteristics, to address the concern that declining occupations would have delivered different outcomes in absence of occupational decline.

To learn more about the mean effect of occupational decline on workers in declining occupations, we consider how the underlying processes affect workers whose occupations do not decline. Workers in non-declining occupations are likely to gain, at least on average, through two channels: first directly, as demand increases for occupations that drive technological change; and second indirectly, as rising incomes increase demand more broadly. Some workers in non-declining occupations may still lose from technological change if their occupations experience small falls in demand, a possibility we address by considering different cutoffs for declining and non-declining occupations. Overall, we expect technological change to benefit the average worker (Caselli and Manning 2019; Humlum 2021). During the period of our study, Swedish workers' incomes rose substantially (for example, Graetz 2020, reports that the real growth of median wages in Sweden was around 1.8% per year). Moreover, mean wages (and earnings) in Sweden increased over the period we study not only in the aggregate, but also for all subgroups formed by age-by-gender-by-education cells. In sum, therefore, our regression estimates likely provide an upper bound on the magnitude of the losses incurred by workers whose occupations decline. Across a range of specifications and robustness checks, we find that this bound is around 2%-5% of earnings from 1985-2013. This indicates that at least on average, Swedish workers in declining occupations were able to avoid large income losses.

As our main measure of occupational decline, we use the US Occupational Outlook Handbooks (Bureau of Labor Statistics 1986; 2018, henceforth OOH), which allows us to identify which occupations declined in the US since the mid-1980s; to check whether occupational declines had likely technology drivers; and to gauge expectations of employment growth at the time. For reasons that we discuss below, our baseline definition of occupational decline requires that employment contracted by at least 25% since the mid-1980s, though we also explore many alternative definitions, including declines using a range of thresholds as well as measure computed using only the Swedish data. We match the occupational information from the OOH to individual-level panel data on the entire Swedish population. Thus, we utilise the best aspects of both countries' data: the US data allow us to characterise occupational employment growth and control for anticipated changes in demand, while the Swedish data let us follow individuals who differ in their exposure to occupational declines, but were otherwise very similar.

Focusing on cohorts that were in prime working age from the mid-1980s till the mid-2010s, we study how cumulative long-run outcomes (such as earnings and employment) differ for those who in 1985 worked in occupations that subsequently declined. We control for the initial selection of workers into declining occupations by gender, age, education, income, and location in 1985. We show that conditional on these controls, those in occupations that subsequently declined had similar cognitive and non-cognitive skills and parental education and earnings, and similar pre-1985 earnings, as other workers. In some specifications we add other controls, including measures of occupation-varying life-cycle profiles and predictors of occupational employment growth, as well as broad occupation and industry dummies. We show using pre-period data that our rich set of controls plausibly addresses the occupation-level differences in outcomes in absence of occupational decline.

We confirm that both our OOH-based measure of occupational decline and the predicted changes in US employment correlate strongly with the employment changes in Sweden. Specifically, Swedish workers who started out in occupations that subsequently declined were exposed to employment growth that was 20-40 log points lower than in non-declining occupations. We find that relative to workers with similar characteristics in non-declining, those in declining occupations lost about 5% of mean cumulative pre-tax earnings and 2% of mean cumulative employment. And compared to similar workers in similar occupations and industries, the cumulative earnings losses were only around 2%, and the cumulative employment losses were around 1%. The implied elasticity of relative employment losses with respect to occupational employment growth ranges from 0.04-0.05, and that for earnings losses ranges from 0.08-0.13.[1,2]

We also find that those in declining occupations were significantly less likely to still work in their 1985 occupation in 2013, and the implied elasticity of remaining with respect to employment growth ranges from 0.71-0.95. If occupational demand curves slope downward, such a strong mobility response likely mitigated the earnings losses for those who remained in declining occupations.

While mean relative earnings losses from occupational decline were around 2%-5%, those in the bottom tercile of their occupation's earnings distribution in 1985 suffered larger relative losses, amounting to 8%-11%. Those at the bottom (and possibly also the top) of their occupation's earnings distribution were also less likely to remain in their starting occupation than the median worker.

We further find that occupational decline increased the cumulative time spent in unemployment (accounting for roughly a third of lost employment) and retraining (accounting for just under ten percent of lost employment). Moreover, occupational decline led to slightly earlier retirement among middle-aged (in 1985) workers. While most of our analysis focuses on overall occupational decline, we also investigate variation stemming from technological replacement using three distinct methods, none of which is mechanically related to the other. We find that all three measures of technological replacement are associated with employment and earnings losses for individual workers that are broadly similar to those in our main estimates.

To frame our empirical analysis of the consequences of occupational decline, we construct a qualitative Roy (1951) model with occupational demand shocks. As discussed above, we find that the largest earnings losses from occupational decline in Sweden are incurred by those who earned the least within their initial occupations. This finding is inconsistent with the frictionless Roy model, but it is consistent with a version where occupational switching costs decline in the workers' ability in the destination occupation. Moreover, our empirical analysis sheds light on the nature of the occupational switching costs,

1. Our paper focuses on changes that result in large-scale occupational declines, rather than task-replacing technologies that change the work done within occupations. Nevertheless, we note that workers performing similar tasks may be exposed to similar demand changes. Indeed, we find that the difference in occupational employment growth between declining and non-declining occupations decreases in magnitude when we control for broad occupational categories. This is reflected in the implied elasticities. There may also be spillovers as workers move from declining occupations into other occupations requiring similar tasks. In a robustness check, we find that spillovers matter to an extent, but do not overturn our main conclusions.

2. We estimate mean earnings and employment losses from occupational decline that are similar in magnitude or possibly even smaller using micro data from the US (National Longitudinal Survey of Youth 1979); the NLSY estimates are, however, noisier than those using Swedish data. Nevertheless, these findings suggest that our estimates of losses from occupational decline may generalise to settings beyond Sweden. See the earlier working paper version for details (Edin et al. 2019).

as almost half of the employment losses we estimate are accounted for by a combination of increased retraining and especially unemployment.

Our model can account for additional empirical findings when we also allow for worker displacement. In this case, those with lower initial within-occupation earnings rank suffer larger earnings losses from occupational decline; switchers' earnings losses may be larger than those of stayers; and switching probabilities are U-shaped in initial earnings, whereby low-earning workers switch occupations if displaced, while high-earning workers switch regardless of displacement when faced with occupational decline.

Taken together, our results suggest that most workers have coped well with occupational decline, in part through successful occupational switching, which is an encouraging message for workers facing the risk of technological displacement today. One exception to this generally positive finding is that low-ranked workers suffer larger losses from occupational decline.

Occupational decline is a salient feature of the evolution of labour markets (Goldin 2000). But despite its importance, past work provides relatively little evidence about its consequences for individuals' careers. While there is a large literature on the costs of occupational mobility, we are not aware of previous estimates of workers' earnings losses from negative occupation-level demand changes.[3]

Our paper is distinct from panel studies of workers who differ in the routineness of their jobs.[4] A key difference is that we can compare similar workers, even doing similar work, with different exposure to occupational decline.[5]

Our paper also differs conceptually from studies of mass layoffs. Occupational decline can sometimes be managed through retirement and reduced hiring, allowing workers to change jobs without leaving employment; and occupational decline need not entail severe spillovers for local labour markets, unlike mass layoffs (Gathmann, Helm, and Schönberg 2018). While magnitude comparisons across studies should be interpreted with caution, the mean loss that we find from occupational decline is generally lower than the loss from mass layoffs.[6] Finally, our paper also differs from studies of trade shocks, which affect import-competing firms and industries, while the changes we study typically affect individual workers within firms.

Our paper is also related to Dauth, Findeisen, Suedekum, and Woessner (2021), who explore how workers fare who are exposed to industrial robots; and to Battisti, Dustmann, and Schönberg (2017), who investigate how firm-level technological and organisational change affects workers' careers. Our paper differs by exploring the consequences of a broader set of changes in occupational employment. Furthermore, our paper is related to the literature on possible future displacement due to technological

3. Cortes and Gallipoli (2017), Kambourov and Manovskii (2009), Pavan (2011), and Sullivan (2010) estimate the human capital losses associated with switching occupations. An older literature, including Neal (1995) and Parent (2000) studies the cost of moving across industries, while in other related work Gathmann and Schönberg (2010) and Poletaev and Robinson (2008) focus on task-specific human capital. Changes in the task content of existing occupations (for instance Spitz-Oener 2006), while also potentially relevant, are outside the scope of our study due to data limitations.

4. See for example Cortes (2016), Autor and Dorn (2009), and Bachmann, Cim, and Green (2019).

5. Also related is independent work by Schmillen (2018), who studies employment shocks faced by German apprentices, although our paper differs in its research question, econometric inference, and outcomes.

6. Studies of mass layoffs in Sweden find losses of 4%-6% of annual earnings in the 5-10 years following displacement (Eliason and Storrie 2006; OECD 2015). In the US losses from displacement are generally larger and range from 7%-14% of earnings (Davis and Von Wachter 2011), or possibly even higher for workers who were highly attached to their firms (Jacobson, LaLonde, and Sullivan 1993). Galaasen and Kostøl (2018) and Bana (2021) explore how mass layoffs' effects differ for occupations facing negative demand shocks, but their focus is still on mass layoffs

changes. Forecasts of occupational displacement range from almost 50% (Frey and Osborne 2017) to around 10% (Arntz, Gregory, and Zierahn 2017, who obtain a lower estimate by taking into account within-occupation heterogeneity in tasks). At the same time, Bessen (2016) concludes that technology has, at least so far, not been a net destroyer of jobs. Even if this benign aggregate trend continues, however, some occupations may be replaced by technology, and our study offers a way to assess the losses from occupational displacement.

We conclude the introduction with brief remarks on the setting we study. Sweden's economy and labour market institutions constitute the backdrop to most of our empirical analysis. During the period of our study, the Swedish economy experienced a deep recession in the early 1990s and a milder one in 2008 (Lindbeck 1997; Gottfries 2018), and we find that earnings losses in declining occupations were worse during those recessions. Wage inequality in Sweden increased during the 1980s and 1990s and remained relatively stable thereafter (Skans, Edin, and Holmlund 2009). Swedish labour market institutions have been characterised by strong labour unions and substantial public spending on labour market policies. Unions have generally embraced technological changes to promote productivity and wage gains, while expecting that active labour market policy will help displaced workers find work (Edin and Holmlund 1995). There is, indeed, some evidence that Sweden's occupational retraining programs raise earnings (Vikström and Berg 2017), so they may have contributed to the modest losses from occupational decline that we find.[7] At the same time, our finding of noisier but similarly modest mean earnings and employment losses from occupational decline in the US (Edin et al. 2019), suggests that workers find ways to mitigate losses from occupational decline even in other settings.

The remainder of our paper is organised as follows. Section 3.2 presents our model, Sections 3.3 and 3.4 discuss our data and empirical strategy, respectively, Section 3.5 presents our results, and Section 3.6 concludes.

3.2 Occupational decline in a Roy model

This section presents a simple model to help us frame our empirical investigation. We consider two occupations, one of which is hit by a negative demand shock, and we qualitatively study the resulting sorting of workers and the costs they incur. We investigate how workers' likelihood of leaving the affected occupation, and their earnings losses, depend on their initial earnings. Starting from a standard frictionless Roy (1951) model, we successively introduce positive and potentially heterogeneous costs of switching occupation; as well as the possibility that workers are displaced from their jobs and incur a cost to find a new job even when remaining in their initial occupation. Finally, we consider how workers' sorting differs when the negative demand shock is anticipated in advance. A complete, self-contained exposition of the model is given in the online appendix. Here we only summarise the main elements.[8]

7. Another feature of Swedish labour market institutions are so-called employment security agreements reached between labour unions and business associations, and administered by works councils. These agreements stipulate counselling of laid-off workers to minimise the duration of their unemployment. We do not consider these agreements important in driving our results because, first, private sector blue-collar workers were only covered from 2004 onwards, and second, a careful evaluation of these agreements does not find strong support for positive treatment effects (Andersson 2017).

8. Gola (2021) provides a different but related theoretical analysis of technological change in a two-sector model.

We consider a competitive economy with a continuum of individuals indexed by i who live for two periods $t \in \{1,2\}$ and each supplies a unit of labour inelastically each period. There are two occupations indexed by $k \in \{A,B\}$ for the workers to choose from. Workers' per-period log earnings are given by $y_{ikt} = \pi_{kt} + \alpha_{ik} - c_{ikt}$, where π_{kt} is the time-varying and stochastic (log) price of a unit of output in occupation k, α_{ik} is the time-invariant (log) amount of output that worker i produces in occupation k, and $c_{ikt} \geq 0$ is a time cost related to occupational switching, which we discuss below.[9] There are no saving opportunities and earnings are consumed immediately. We define the life-time expected utility function as $\mathbb{E}[y_{ik1} + \beta y_{ik2}]$, where $\beta > 0$ is a discount factor. In each period, workers choose the occupation that maximises their expected utility. As a normalisation, we assume that workers always choose occupation A if indifferent. Since we focus our analysis on relative wages, we define $\widetilde{\pi}_t \equiv \pi_{Bt} - \pi_{At}$ and assume for simplicity that $\widetilde{\pi}_1 = 0$. As we discuss in Section 3.4.1, our empirical analysis proceeds in two steps. First, we characterise conditions for identifying the effect of occupational decline on those in declining occupations relative to others. And second, we discuss the additional assumptions required to identify a bound on the absolute losses of those in declining occupations. The model is focused on the first part. Prices are determined in equilibrium by supply and demand. However, here we take them as given, and analyse the consequences of a change to prices occurring in period 2 for occupational sorting and earnings. Note that the second period may be interpreted as all periods following this change, so β could be larger than one. For simplicity, we assume that α_{iA} and α_{iB} are independent and both uniformly distributed between zero and some finite but possibly large number $\overline{\alpha}$. We explain in the online appendix that our main results are robust to alternative assumptions about the joint skills distribution.

In period 2, there is a negative demand shock to occupation A such that $\pi_{A2} - \pi_{A1} = -d$ and $\widetilde{\pi}_2 = d, d > 0$. This may be due to labour-replacing technology becoming available, or cheaper, in occupation A. We are interested in the consequences of the shock for the earnings of workers who start out in occupation A, under various assumptions about switching costs and anticipation of the price change. Formally, let $l_i \equiv \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 1]$ be the expected earnings loss in period 2 that results from the shock, conditional on worker i starting out in occupation A, and conditional on her ability (and hence earnings rank) α_{iA}, where the occurrence of the shock is indicated by $D_A \in \{0,1\}$. Similarly, l_i^{switch} and l_i^{stay} denote the earnings losses further conditioned on leaving and staying, respectively, and p_i is the probability of switching.[10] The overall loss is given by

$$l_i = l_i^{\text{stay}} - p_i \left(l_i^{\text{stay}} - l_i^{\text{switch}} \right). \tag{3.1}$$

As long as there is no displacement then $l_i^{\text{stay}} = d$ and by revealed preference $l_i^{\text{switch}} \leq d$, so that $l_i \leq d$.

9. The time cost may reflect search or retraining (or both); we assume throughout that a worker's wage equals the value of her marginal product, $e^{\pi_{kt}+\alpha_{ik}}$. We thus abstract from any job-level rents that may arise in the presence of search frictions.

10. Formally,

$$\begin{aligned}
l_i &\equiv l_i(\alpha_{iA}, d) \equiv \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 1],\\
l_i^{\text{switch}} &\equiv l_i^{\text{switch}}(\alpha_{iA}, d) \equiv \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, k_{i2} = B, D_A = 1],\\
l_i^{\text{stay}} &\equiv l_i^{\text{stay}}(\alpha_{iA}, d) \equiv \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, k_{i2} = A, D_A = 1],\\
p_i &\equiv p_i(\alpha_{iA}, d) \equiv \mathbb{P}(k_{i2} = B|k_{i1} = A, \alpha_{iA}, D_A = 1).
\end{aligned}$$

Thus, switching enables workers to mitigate the losses from occupational decline. In the online appendix we show that in each version of our model, $\frac{\partial p_i}{\partial d} \geq 0$, $\frac{\partial l_i}{\partial d} \geq 0$ (with strict inequalities for some i): the larger the drop in demand, the more workers switch, and the higher are earnings losses. Furthermore, $\frac{\partial l_i}{\partial \alpha_{iA}} = -\frac{\partial p_i}{\partial \alpha_{iA}} \left(l_i^{\text{stay}} - l_i^{\text{switch}} \right) + p_i \frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}}$. In other words, losses decrease with initial within-occupation earnings rank if the switching probability is increasing and the loss of switchers decreasing in initial earnings rank, $\frac{\partial p_i}{\partial \alpha_{iA}} > 0$ and $\frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} < 0$.

We start with the simplest case, where occupational prices π_{kt} are revealed at the start of each period and there are no switching costs. Hence, occupational choice is a sequence of static decisions that can be analysed in isolation. Panel (a) of Figure 3.1 illustrates occupational choices in the two periods as a function of workers' skills. The set of workers who start out in occupation A but then switch to B is indicated by the blue area in the figure. Given uniformly distributed skills, the figure shows that $\frac{\partial p_i}{\partial \alpha_{iA}} \leq 0$. We show in the online appendix that also $\frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} \geq 0$, and that $\frac{\partial l_i}{\partial \alpha_{iA}} > 0$: mean losses from occupational decline increase with initial earnings.

Figure 3.1 about here

To understand the intuition for these results, call occupation A "typist" and occupation B "cashier", where typists suffer a negative demand shock. The worst typists could only become the worst cashiers, otherwise they would have chosen to be cashiers in period 1. But the best typists can at most become the best cashiers, and in general they will not all be the best cashiers. Therefore, the best typists are less able to mitigate their earnings losses by becoming cashiers, and they suffer larger losses than the worst typists. This argument suggests that switching probabilities are decreasing and earnings losses are increasing in ability under a large set of alternative assumptions on the skill distributions.

Next, we assume there is a constant switching cost $c \in (0, d)$ for moving between occupations. Occupational choice is no longer a period-by-period decision. Instead, workers choose in period 1 the occupation with the highest expected present discounted value of log earnings, net of switching costs. Let us assume that occupational log prices follow a random walk, $\mathbb{E}[\widetilde{\pi}_2] = \widetilde{\pi}_1 = 0$, where the last equality is due to our earlier simplifying assumption.[11] Panel (b) of Figure 3.1 shows that occupational choices are qualitatively similar to the baseline model, except that the blue region marking the workers who switch is smaller than in panel (a). Again we have $\frac{\partial l_i}{\partial \alpha_{iA}} > 0$.

Instead of a constant switching cost, let us now assume that the cost for moving from A to B equals $C - \alpha_{iB}$, with $C > \overline{\alpha}$ (symmetrically, the cost of moving from B to A equals $C - \alpha_{iA}$). This structure of switching costs captures in a reduced form way the frictions that occupational moves may entail: for example, job search may take time, and those more able in the new occupation may find a job more quickly. We continue to assume that occupational log prices follow a random walk. Panel (c) of Figure 3.1 shows that low-ability workers do not leave occupation A, and among high-ability workers, $\frac{\partial p_i}{\partial \alpha_{iA}} > 0$. We show in the online appendix that $\frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} < 0$ (taking into account earnings losses due to the time cost

11. The random walk assumption is consistent with our scenario of occupation A experiencing an adverse shock in period 2. Since occupations are completely symmetric ex-ante, this is without loss of generality. Instead of the random walk assumption we could impose that demand changes are somehow otherwise perfectly unforeseen, for instance due to adaptive expectations.

of switching), so that $\frac{\partial l_i}{\partial \alpha_{iA}} \leq 0$: mean losses from occupational decline (weakly) decrease with initial earnings.

In terms of the example above, in this case the worst typists do not switch, because their initial choice of occupation A reveals not only low earnings potential in occupation B but also a large switching cost. Among the best typists, however, many possess substantial earnings potential as cashiers, as well as low switching costs. Therefore, the best typists are on average better able to mitigate their earnings losses by becoming cashiers, and hence the earnings losses from the demand shock are smaller for the best typists than for the worst typists.[12]

Building on the previous case, we now explore a version of the model that includes involuntary job displacement. Suppose that workers experience job displacement with some probability that is independent of skill, and incur a time cost $C - \alpha_{ik}$ to find a job in occupation k, be it the starting occupation or a different one. Here we have in mind exogenous job losses, for instance due to plant closure, which are a standard feature of search models (see for instance Pissarides 2000). There is a fraction of high-ability workers who switch occupation regardless of displacement. In addition, now a fraction of low-ability workers also switch, but only if they are displaced. This is illustrated by the yellow area in panel (c') of Figure 3.1. Moreover, the earnings losses experienced by these displaced movers are larger than those of non-displaced, comparable stayers. This is by revealed preference: a worker in the yellow region prefers to remain if not displaced, so her non-displaced counterpart (with the same period-1 earnings) necessarily incurs a lower earnings loss. We show in the online appendix that $\frac{\partial l_i}{\partial \alpha_{iA}} \leq 0$, as before. Unlike in the case without displacement, however, p_i is U-shaped in initial earnings. This is because the probability of a displacement-induced switch is decreasing, and that of a voluntary one is increasing in initial earnings. The earnings loss l_i is again decreasing in initial earnings, as the costs of moving jobs—both within and across occupations—decrease with initial earnings.[13]

As a final variation on our model, we consider a case where period-2 prices are revealed to be $\widetilde{\pi}_2 = d$ at the start of period 1. In the presence of switching costs, some workers that would otherwise have chosen occupation A in period 1 instead start out in occupation B. This means that the fraction of workers switching after period 1 is smaller, and it could even be zero if switching costs are large. Since there is less switching, earnings losses are larger than in the case of unanticipated shocks, for a given d.

We conclude this section by summarising the main results from our model. The baseline frictionless model makes three predictions: the probability of leaving a declining occupation is decreasing in initial earnings; earnings losses due to occupational decline are increasing in initial earnings; and earnings losses of those who leave a declining occupation are less than the losses of those who remain. Anticipating that these predictions are inconsistent with our empirical findings, we consider how occupational switching costs can reconcile our results. Introducing an occupational switching cost that is decreasing

12. While our model excludes occupation-specific human capital, it does allow us to think about some of its potential implications. For example, if all workers accumulate occupation-specific human capital additively (in logarithms) the effects are similar to adding constant switching costs, since switching means foregoing this capital. And if the accumulation of occupation-specific human capital is faster for those with higher ability in the occupations they initially select, they become less mobile, in contrast to the case of heterogeneous switching costs discussed above. Either way, adding occupation-specific human capital does not help to rationalise our empirical findings.

13. We have also analysed displacement under constant switching costs, that is, when workers incur a time cost $\widehat{c} > 0$ to find a new job in A, or a cost c to find a job in B. This case is illustrated by panel (b') of Figure 3.1, and details are given in the online appendix.

in the worker's earnings in the destination occupation, leads to a positive relationship between switching probabilities and initial earnings, and a negative relationship between earnings losses and initial earnings. Allowing for displacement, together with a cost of switching jobs within an occupation, implies that switchers' earnings losses may be larger than those of stayers. Moreover, displacement can cause switching probabilities to be U-shaped in initial earnings, whereby low-earning workers switch involuntarily if displaced, while high-earning workers switch voluntarily regardless of displacement.[14]

3.3 Data

Our main empirical analysis uses individual-level longitudinal administrative data covering the entire population of Sweden for several decades, and various editions of the Occupational Outlook Handbook (OOH) published by the Bureau of Labor Statistics (BLS). Here we discuss key elements of the data we use, and leave many of the details to the online appendix.

3.3.1 Data sources

Our primary sources for measuring occupational decline are the 1986-87 and the 2018-19 editions of the Occupational Outlook Handbook (Bureau of Labor Statistics 1986, 2018). The OOH describes the nature of work, the number of jobs, and the projected employment growth for hundreds of occupations. For a subset of these occupations, more details are reported, including (among much else) data on whether technology is expected to affect—or has already affected—the occupation in question, and if so in what way. In the 1986-87 edition, 401 occupations are described, covering about 80% of US employment. Detailed information is available for 196 of these occupations, covering about 60% of employment.[15]

Our main outcomes of interest come from Swedish micro data. We obtain basic demographic (year of birth, gender, education, and county of residence) and labour market (employment status, annual earnings, and industry) variables from the Integrated Database for Labour Market Research (LISA), a collection of administrative registers. For 1985-2013, LISA contains one observation per year for every individual aged 16-64 living in Sweden. Key variables, such as employment status and industry (as well as county of residence) are measured each November. We also use individual data from the Swedish Public Employment Service (PES), which contain information on the number of days registered as unemployed and number of days spent in retraining programs administered by the PES, for all individuals ever registered with the PES from 1992-2013.

To assess balance between treatment and control groups in terms of pre-determined characteristics, we use information on cognitive skills (an IQ-type measure) and non-cognitive skills (capturing psy-

14. An alternative model of occupational decline is the hierarchical Roy model of Cortes (2016). There are three occupations that differ by skill intensity—there is only one dimension of skill—and the declining occupation is assumed to be the middle-skilled one. Among middle-skilled workers, both the lowest and highest paid leave the occupation, while the medium paid workers stay. As in any frictionless Roy model, stayers incur the largest earnings losses when an occupation declines. Therefore, the model of Cortes (2016) cannot explain our finding that the lowest-paid within the occupation incur the largest losses. And naturally, that model cannot speak to our findings about unemployment and retraining.

15. The number of distinct occupations in the OOH, as well as the number of occupations covered in detail, increased over time, so our crosswalk from the 1986-87 to the 2018-19 OOH is mostly, though not always, one-to-many.

chological traits such as the ability to cope with stress) from the military enlistment. These data are described in detail by Lindqvist and Vestman (2011). We also use information on parents' education and income from the 1985 version of LISA.

Our data on workers' occupations come from the population censuses, which were conducted every five years from 1960-1990, and from the Wage Structure Statistics (WSS) for the years 1996-2013.[16] The WSS contains the population of public sector workers and a sample of about 50% of private sector workers. We apply sampling weights when working with the occupation variable from the WSS.

A useful feature of our data is that in the 1985 and 1990 censuses, workers' occupation is coded using a 5-digit classification, YRKE5, containing about 1,400 distinct occupations. This allows us to accurately merge occupation-level information from the US, as we describe below. Unfortunately, such detailed occupation codes are not available after 1990. From 1996-2013, a 3-digit classification containing 172 distinct codes, SSYK96, is available in the WSS. This classification is different from YRKE5, and the cross-walk between YRKE5 and SSYK96 likely introduces measurement error in workers' occupations after 1990. This limits our analysis of occupational employment shifts and individual workers' occupational mobility during 1985-2013.

Finally, we use information from the 1960 census, which allows us to calculate prior occupational employment changes at the 3-digit level using the YRKE3 classification, a coarser version of YRKE5 (there are 229 distinct codes that cover the period 1960-85).

3.3.2 Construction of key variables

To construct our measure of occupational decline we begin with the OOH data. Mapping occupations across the 1986-87 and 2018-19 editions of the OOH, we calculate the percentage growth in employment 1984-2016.[17] If, after a careful search, a 1986-87 occupation has no counterpart in the 2018-19 edition, we classify it as having vanished, and assign a percentage growth of -100.[18] While few occupations actually disappeared, some occupations declined sharply, including both white-collar (typists, drafters, and telephone operators), and blue-collar (precision assemblers, welders, and butchers) jobs.

We also record the projected employment growth of each occupation from the 1986-87 OOH.[19] The BLS constructs these predictions using a careful and lengthy procedure.[20] In the 1986-87 OOH, forecasts were reported in categories: "declining", "little or no change", "increasing slower than average", "increasing about as fast as average", and "increasing faster than average". We create a cardinal predicted growth index assigning these categories the numbers 1-5 (where higher numbers correspond to more positive predicted employment changes). We report results both from using this index and using the categorical outlook variable.

In order to merge the OOH-based variables to the Swedish data, we map the 401 1986-87 OOH oc-

16. We also use individual-level earnings data for 1975 and 1980 from the population censuses.
17. The 1986-87 OOH reports employment for 1984, while the 2018-19 edition reports 2016 employment.
18. Between the 1986-87 and 2018-19 editions of the OOH, some occupations were split or merged, which we take into account when computing the percentage growth. See the online appendix for details.
19. We use US rather than Swedish employment forecasts since the Swedish forecasts from the relevant time period are only at the one-digit occupation level (Statistics Sweden 1990).
20. Veneri (1997) evaluates the ex-post accuracy of the projections used in the 1986-87 OOH, and concludes that it correctly foresaw most occupational trends, although there were non-trivial cases of error.

cupations to the 1,396 5-digit Swedish occupation codes available in the 1985 census. We successfully map 379 US occupations to 1,094 Swedish occupations—we are able to find corresponding US occupations for 91% of Swedish workers in 1985. We map percentage changes in US employment 1984-2016, as well as 1986-87 OOH predictions (categorical and index), to Swedish 5-digit occupations using our crosswalk, applying weights (OOH 1984 employment shares) in the case of many-to-one matches.

We define a Swedish 5-digit occupation as declining if the weighted employment growth of its corresponding OOH occupations is negative and larger (in absolute magnitude) than 25%. We regard this as a sensible threshold: smaller observed declines may result from measurement error from matching OOH occupations over time. At the same time, we report robustness checks using a number of alternative thresholds. We also use information from the OOH to determine whether technology likely played a role in the decline, as we further explain in the online appendix. In 1985, 13% of Swedish employees worked in subsequently declining occupations, and 8% worked in subsequently declining occupations where the decline is linked to technological change. Examples of sharply declining occupations in Sweden include both blue-collar (vehicle fitter, assembler of metal products, machine fitter, and clothing seamstress) and white-collar occupations (office telephone operator, data reader, data machinist, and typist).[21]

We also classify occupations as having been susceptible to replacement by specific technologies. Unlike the declines linked to technological change, these occupations were categorised without relying on actual employment changes in the US (nor Sweden). We identify relevant technologies using two approaches: a 'manual' one and an 'algorithmic' one. For the manual approach, we consider whether we know of a technology that replaced all or nearly all of tasks in the occupation. For the algorithmic approach, we use a pre-specified Google search query to identify mentions of technology replacing workers in the occupation. An occupation is considered to have been replaced if the technology identified was commercially viable during the period we study. In both cases, technological replacement is strongly positively correlated with occupational decline. In 1985, 5.4 (3.7)% of Swedish employees worked in occupations that were replaced, as classified using the algorithmic (manual) approach.

We construct several left-hand side variables that characterise workers' career outcomes spanning the years 1986-2013, that is, starting with the first year after we measure treatment and ending with the last year available in our data. We start by summing up years observed as employed and real annual labour earnings, obtaining the variables cumulative years employed and cumulative earnings. Following Autor, Dorn, Hanson, and Song (2014), we measure normalised cumulative earnings, which is the ratio of cumulative earnings to predicted initial earnings.[22] We consider further earnings measures—such as rank, discounted cumulative earnings, and earnings growth—in robustness checks.

Our measure of long-run occupational mobility is a dummy variable that equals one if the individual

21. This list includes occupations with at least 5,000 workers in 1985; whose decline measure fell by 50% or more; and with distinct names, as opposed to "Other within [a broader category]".

22. The prediction comes from a regression of log earnings on a quartic in age and dummies for gender, county, and seven education categories, run separately for each 3-digit SSYK96 occupation in 1985. We divide by predicted rather than actual initial earnings to eliminate transitory earnings variation, which would introduce an important role for mean reversion into the distribution of normalised cumulative earnings. Autor, Dorn, Hanson, and Song (2014) divide cumulative earnings by earnings averaged across four pre-treatment years for the same reason. Since we do not have annual earnings information prior to 1985, we normalise by predicted earnings instead.

worked in the same 3-digit SSYK96 occupation in 2013 as 1985. It equals zero if the individual works in a different occupation or is not employed.[23] Using the PES data, we calculate cumulative days spent unemployed and cumulative days spent in retraining during 1992-2013. We define dummy variables for ever unemployed and ever having participated in retraining. Finally, we calculate the retirement age, where we define retirement as a continuous spell of zero annual earnings up to and including age 64.[24]

3.3.3 Sample restrictions

Our starting sample contains all individuals born between 1921-1969 and hence aged 16-64 (at some point) in 1985; who were employed in November 1985; whose annual earnings in 1985 were no less than the "base amount" (Swedish: *basbelopp*) specified by the social security administration; and for whom we have the relevant demographic and labour market information.[25] There are 3,061,051 individuals fulfilling the above criteria.[26] Our *baseline sample* further restricts the sample to those aged 25-36 in 1985. We drop younger workers, who are less likely to have settled on an occupation. And we drop middle-aged and older workers from our baseline sample because we want to focus on workers who did not reach retirement age by 2013, the end of our period of study, in our main analysis. We analyse these older workers separately.

3.4 Empirical strategy

3.4.1 The estimating equations and their interpretation

Our objective is to estimate the consequences of occupational decline for individual workers' careers. To fix ideas we consider occupational decline brought about by technological change, and later consider other potential drivers of occupational decline.

Consider a regression of cumulative career outcomes—such as cumulative years employed, or cumulative earnings—on an indicator for working in 1985 in occupations that subsequently declined, conditional on a set of controls. The probability limit of the regression coefficient on the declining indicator can be expressed as a difference in conditional means, which in turn can be decomposed into the difference between a treatment effect on workers in declining occupations and a treatment effect on

23. Our measure of occupational mobility does not capture any temporary exits during the intervening years if workers returned to their initial occupation. A limitation of our data is that they are not conducive to studying high-frequency occupational mobility: During the years 1986-1989 and 1991-1995, we do not observe workers' occupation. And during 1996-2004, the SSYK96 variable contains substantially fewer distinct codes than from 2005 onwards.

24. The LISA database includes individuals older than 64 only during later years. Since we do not consistently observe individuals beyond age 64, we assume for all years that individuals aged 65 or older have retired.

25. The base amount is used as an accounting unit when calculating benefits, and it is typically equal to about three months' worth of full-time work at the median wage. As we do not observe hours worked or fulltime status, we use the base amount to exclude individuals with little labour market attachment.

26. There were 5,281,382 individuals aged 16-64 in Sweden in 1985. Of those, 4,186,512 were employed in November 1985, and among them, 3,648,034 earned no less than the base amount during 1985. The reduction to 3,061,051 is due to missing education, industry, or occupation information, including cases where YRKE5 occupations do not have matches in the OOH. Table 3.8 shows that dropping observations with missing education, occupation, and industry has very little effect on the sample composition in terms of gender, age, and base earnings.

workers in non-declining occupations, and selection bias:

$$\underbrace{\mathbb{E}[y_{i2}|k_{i1}\in A, D_A=1, D_B=0, \mathbf{x}_{i1}]-\mathbb{E}[y_{i2}|k_{i1}\in B, D_A=1, D_B=0, \mathbf{x}_{i1}]}_{\text{Observed difference in means}} =$$

$$\underbrace{\mathbb{E}[y_{i2}|k_{i1}\in A, D_A=1, D_B=0, \mathbf{x}_{i1}]-\mathbb{E}[y_{i2}|k_{i1}\in A, D_A=0, D_B=0, \mathbf{x}_{i1}]}_{\text{Effect of occupational decline on } A \text{ workers}}$$

$$-\underbrace{(\mathbb{E}[y_{i2}|k_{i1}\in B, D_A=1, D_B=0, \mathbf{x}_{i1}]-\mathbb{E}[y_{i2}|k_{i1}\in B, D_A=0, D_B=0, \mathbf{x}_{i1}])}_{\text{Effect of occupational decline on } B \text{ workers}}$$

$$+\underbrace{\mathbb{E}[y_{i2}|k_{i1}\in A, D_A=0, D_B=0, \mathbf{x}_{i1}]-\mathbb{E}[y_{i2}|k_{i1}\in B, D_A=0, D_B=0, \mathbf{x}_{i1}]}_{\text{Selection bias}} . \tag{3.2}$$

Here, y_{i2} is the outcome of interest, such as cumulative earnings of individual i (who is employed in occupation k_{i1} in 1985) in period 2 (1986 through 2013). Our notation separates declining occupations (A) from non-declining ones (B). D is an indicator for occupational decline, which allows us to consider the hypothetical situation where occupational decline did not take place. We motivate our use of an indicator for occupational decline in Section 3.4.2 below. $\mathbf{x}_{i1}$ is a set of (yet unspecified) controls, which we also revisit below.

In our model, the selection term equals zero due to the symmetry assumptions we make, and there is no need for controls. In practice, the identifying assumption for the regressions without controls is too strong, because the selection term reflects both differences in individuals sorting across occupations as well as differences between occupations A and B even in absence of occupation decline. Our empirical strategy aims to mitigate both of these types of selection concerns.

Our first step towards addressing the sorting of individuals is to control for a rich vector of individual characteristics in period 1, that is, in 1985: gender, age, educational attainment and county of residence dummies, and earnings. To address potential concerns about sorting on other dimensions of skill, we investigate whether individuals differ in terms of cognitive and non-cognitive skills from the military enlistment, which are available for men of a subset of cohorts for whom military service was compulsory. We also check whether workers in declining occupations differ in other background characteristics, namely parental education and pre-1985 earnings.

But even if the controls resolve concerns about sorting on individuals' characteristics, the selection term may still be non-zero if earnings growth in occupations A and B would have been different in the absence of occupational decline—that is, if a worker's occupation affects her earnings growth even without occupational decline. To mitigate this concern, we use data from 1985 to estimate expected occupation-specific lifetime earnings profiles, which we add as controls to the regressions.

To further control for differences between declining and non-declining occupations, we use information from the 1986-87 OOH. The BLS authors went to great lengths to accurately forecast occupational employment changes. Once we condition on predicted occupational growth, we likely remove much of the differences between declining and non-declining occupations that are related to workers' sorting in anticipation of future demand. Since the occupational decline and the forecasts that we use rely on US data, we also control for each occupation's level of employment and pre-existing employment growth trends in Sweden. Together, this set of controls removes predictable variation in the declining indicator, and in this sense gets us closer to isolating unanticipated declines. In some specifications we use two additional sets of controls: broad (1-digit) occupation dummies and (2-digit) industry dummies. Adding these controls comes at the cost of reducing the variation in occupational decline, since it only uses variation in occupational decline between very similar occupations.

Addressing the selection concerns allows us to identify the net effect of occupational decline on workers who start out in occupations that subsequently decline, relative to the effect of occupational decline on workers who start out in non-declining occupations (the difference of the two middle terms in equation (3.2)). This difference is equivalent to $-\mathbb{E}[l_i]$ in the model and is interpretable as the effect of occupational decline on inequality between occupations.

To view equation (3.2) as informing us about the effect on earnings of occupational decline on workers who start out in declining occupations, we need to make further assumptions about its effect on workers in non-declining occupations. There are reasons to think that this effect would be at least weakly positive, even if not all non-declining occupations gain. Increased demand due to technological change may raise demand for workers in non-declining occupations and may open new employment opportunities (Autor 2015), and technological change should lead to higher average wages under general conditions (Caselli and Manning 2019). On the other hand, there may be a secondary effect due an inflow of workers from declining occupations to non-declining ones, but its magnitude is likely to be modest (for calculations based on structural models that are consistent with this intuition, see Acemoglu and Restrepo 2022 and Humlum 2021). Finally, we note that real earnings and wages in Sweden have indeed increased substantially during our sample period (Graetz 2020), and the most likely explanation for such growth over long periods of time is a technology-driven increase in productivity. Moreover, mean wages (and earnings) in Sweden increased over the period we study not only in the aggregate, but also for all subgroups formed by age-by-gender-by-education cells.[27] Under the assumption that workers in non-declining workers do not, on average, lose from occupational decline, we get an inequality:

$$\underbrace{\mathbb{E}[y_{i2}|k_{i1} \in A, D_A = 1, D_B = 0, \mathbf{x}_{i1}] - \mathbb{E}[y_{i2}|k_{i1} \in B, D_A = 1, D_B = 0, \mathbf{x}_{i1}]}_{\text{Observed difference in means}} \leq$$

$$\underbrace{\mathbb{E}[y_{i2}|k_{i1} \in A, D_A = 1, D_B = 0, \mathbf{x}_{i1}] - \mathbb{E}[y_{i2}|k_{i1} \in A, D_A = 0, D_B = 0, \mathbf{x}_{i1}]}_{\text{Effect of occupational decline on } A \text{ workers}} . \tag{3.3}$$

In words, the observed difference in conditional earnings means is a lower bound for the effect of occupational decline on workers in declining occupations. Since, as we discuss below, we find that the observed difference is moderate in magnitude, the loss for those in declining occupations is even more moderate.

27. Specifically, we break the population of workers into three age groups (25-36, 37-48, and 49-60, as we do elsewhere in this paper), by five education groups, and by gender. For all the resulting 30 groups, real earnings and real wage rates increased by at least 1.2$ and 1.4% per year, respectively, from 1985-2013.

To implement the empirical methodology outlined above, our estimating equation takes the form

$$y_{i2} = \beta D_{k_{i1}} + \gamma \mathbf{x}_{i1} + \delta \widetilde{\mathbf{x}}_{k_{i1}} + \varepsilon_i, \tag{3.4}$$

where $D_{k_{i1}}$ is an indicator for working in 1985 an occupation that subsequently declined; $\mathbf{x}_{i1}$ is a vector of individual characteristics, measured in 1985, as discussed above; $\widetilde{\mathbf{x}}_k$ is a vector of occupational characteristics; and ε_i is the error term, which we conservatively cluster by three-digit Swedish occupations.[28]

We provide further evidence that our identification strategy plausibly addresses the selection issues by considering two additional sets of outcomes. First, we examine the earnings of workers in the years before the occupational decline that we study. Second, we study the cumulative earnings during the first few years of our study, when the effect of occupational decline were likely limited.

A different question regarding our approach is whether occupational decline that is specifically linked to labour-replacing technologies has distinct consequences from demand-driven occupational decline in general. To provide evidence on the role of technology, we restrict some of the analysis to occupational declines that are explicitly linked to concrete new technologies, such as personal computers and robots. We also examine measures of technological replacement that are solely based on an occupation's exposure to labour-replacing technologies without incorporating any information on employment changes.

Returning to the potential causes of occupational decline, we note that losses due to increased international trade or offshoring are likely to follow a similar logic to the one outlined above for technological changes, although admittedly there is less evidence on this.

One final step in our empirical analysis is to consider how the costs of occupational decline may fall differently on workers with different initial earnings rank within their occupations, in line with the discussion in the model section.

3.4.2 Rationale for measuring occupational decline using US data

Prior literature has documented that shifts in occupational employment are strongly correlated across countries, see for instance Goos, Manning, and Salomons (2014) documenting job polarisation across European countries, and in particular Adermon and Gustavsson (2015) on job polarisation in Sweden. Here we explain why using measures of occupational decline from the OOH is not only feasible, but also desirable.

We begin by explaining why we prefer this measure of decline to an alternative measure using the SSYK96 codes. First, there are 401 OOH codes compared to just 172 (three-digit) SSYK96 codes, and having more codes affords us more variation from small and declining occupations. For example, it lets us separate typists, whose employment fell sharply, from secretaries, whose employment grew. To use the OOH codes we match them to YRKE5 codes, but since the YRKE5 are more numerous we do not lose much variation. Second, since the SSYK96 codes were introduced from 1996 they reflect a

28. As discussed above, some specifications also control for industry fixed effects.

judgement on an occupation's importance made after the start of the occupational decline that we study. Consequently, SSYK96 are more likely to pool occupations with low employment in 1996 (including declining ones) with non-declining occupations. Because the 2018-19 OOH separately describes even occupations with very low employment, this is less of a problem for our approach. Finally, using occupational declines measured in Sweden as a regressor where the dependent variable is change in earnings creates a problem of simultaneity. This problem is mitigated by using the OOH measure.

At this stage readers may also ask: why do we report reduced form results using the OOH decline measure rather than use it as an instrument for occupational decline measured in Sweden using SSYK96? Our rationale for the reduced form approach is that it preserves much more of the variation that we are interested in, for several reasons. First, as noted above, if we use measures based on SSYK96 codes, we lose much of the variation in occupational decline because of the coarseness of the classifications and the lower likelihood of separating occupations in decline. Second, 2SLS would exacerbate this problem, since it only uses part of the variation in the decline. Finally, as we show below, while we still have power to detect changes in occupational decline in Sweden, once we control for predicted changes we are left with a weak instrument.

Nonetheless, we sometimes report the implied elasticities of our outcome variables with respect to occupational employment growth in Sweden as an additional way of assessing magnitudes. Given the caveats just mentioned, these elasticities are likely upper bounds, as the measured Swedish employment change is biased towards zero due to the absorption of declining occupations into broader categories in the SSYK96 classification.

Still another question is why we focus only on occupational declines instead of using the full variation in OOH occupational change. Again there are several factors that influence our choice. First, declines are interesting from the perspective of their social costs and policy implications. Second, large declines in employment are likely driven by declines in labour demand, and we use several alternative measures of technological replacement to corroborate our findings. Finally, as we explain below, we use different cutoffs in the regressions as well as graphical evidence to show that the costs of occupational change are concentrated among those who experience substantial occupational declines; increases or moderate declines seem to matter little relative to each other. Nevertheless, for completeness we also report estimates using the full variation in occupational changes.

To conclude, we note that in order to better assess the quantitative importance of the estimated earnings losses, we relate them to the estimated impacts on occupational mobility, as well as to Swedish occupational employment growth.

3.5 Empirical analysis

In this section we present the findings from our empirical analysis. First, we quantify occupational decline in Sweden and discuss sorting into declining occupations. Second, we study how employment, earnings, and occupational mobility differed for workers in declining occupations. Third, we investigate how the consequences of occupational decline differed by workers' initial within-occupation earnings rank. Fourth, we explore some of the mechanisms through which occupational decline operates, including unemployment, retraining, and early retirement. Fifth, we examine whether occupational declines with observed links to technology has distinct consequences. At the end of this section, we interpret our findings through the lens of the theoretical model from Section 3.2.

3.5.1 Occupational decline and sorting across occupations

We begin by quantifying workers' exposure to occupational decline. In Table 3.1 we report estimates of equation (3.4), where the dependent variable is log employment change from 1985-2013 in each worker's three-digit occupation. Panel A shows estimates for workers aged 16-64 in 1985, and Panel B focuses on our main sample of workers—those aged 25-36 in 1985. The results, which are similar across panels, suggest that workers in declining occupations are exposed to a log employment change that is about 50 log points lower than for workers in non-declining ones; about 40 log points lower when we compare observationally similar workers; and lower by about 20-30 log points when we also include occupation and industry controls. It is important to keep these results in mind when interpreting our findings from putting individual-level career outcomes on the left-hand side of equation (3.4). While adding more controls reduces the risk of omitted variables bias, the results in Table 3.1 show that this also leaves less variation in exposure to occupational change. We also note, as discussed in Section 3.4.2, that these estimates likely understate the employment decline for 5-digit occupations, which we are unable to measure.[29] Table 3.9 reports similar estimates, aggregated by three-digit occupations and weighted by 1985 Swedish employment shares, using our main sample of workers.

Table 3.1 about here

Having described the extent of occupational declines, we turn to the sorting of individuals in 1985 into subsequently declining occupations. Table 3.2 presents results from regressions of several individual characteristics on an intercept and the declining indicator. The top panel considers the entire working-age population with non-missing demographic and labour market information, and the bottom panel focuses on our main sample. In both cases, the sorting patterns are similar: those in occupations that subsequently declined were of similar age, and more likely to be male, less educated, and more likely to be employed in manufacturing. Coincidentally, the gender gap in earnings is offset by the differences in schooling, and on net, the workers in subsequently declining occupations had similar earnings to others in 1985.

Table 3.2 about here

29. The difference of 50 log points translates into an employment decline of about 18%. Let y_i be the log employment change assigned to each individual based on her 1985 5-digit occupation, and D_i be the declining indicator. From the regression $y_i = \alpha + \beta D_i + \varepsilon_i$ we obtain $\mathbb{E}[\exp\{y_i\}|D_i = 1] = \exp\{\alpha + \beta\}\mathbb{E}[\exp\{\varepsilon_i\}|D_i = 1]$, and plugging in our estimates, $\exp\{0.06 - 0.49\} \times 1.26 = 0.82$. As we argue in the text, this likely understates the actual average decline in the 5-digit occupations.

We next investigate whether there is sorting into declining occupations based on cognitive skills, non-cognitive skills, and parental attributes, and if so, whether any differences in these variables disappear once we control for the individual characteristics described in the previous paragraph. Columns (1) and (4) in panel A of Table 3.3 show that in 1985, the cognitive and non-cognitive skills of men in subsequently declining occupations were lower than those of the other men by about 0.2-0.25 standard deviations. But the table also shows that our set of demographic controls (columns (2) and (5)) and additional controls largely solves the selection problem, as the estimates with controls are small and not statistically significant. Panels B and C repeat the analysis for mothers' and fathers' schooling and earnings in 1985, which have the advantage of being available for women as well as men. The pattern is qualitatively very similar to that of the military skill measures: working in a subsequently declining occupation is associated with lower parental schooling and earnings, but these differences disappear once adding controls. In sum, Table 3.3 suggests that, although there is negative selection into occupations that later decline, most of this selection can be mitigated using suitable controls such as education and 1985 earnings. To the extent that minor negative selection persists in the regressions with individual controls, these regressions may slightly overstate the negative effect of occupational decline on workers, but this is not a concern once we add all the controls.

Table 3.3 about here

As a final check for sorting into declining occupations, we investigate earnings in 1980 for the older cohorts in our baseline sample (individuals aged 25-31 in 1980).[30] Again we find that conditional on individual-level characteristics, there are essentially no differences in prior earnings, as seen in Figure 3.4. Taken together, the results in this section suggest that concerns about sorting into declining occupations are largely alleviated when we include suitable controls.

3.5.2 Main results on employment, earnings, and occupational mobility

Table 3.4 reports results from estimating equation (3.4) using our main sample of workers aged 25-36 in 1985. Panel A shows that workers in declining occupations spent about nine fewer months (0.73 fewer years) in employment from 1986-2013 (column (1)). Once we add individual controls, this estimate reduces to about six months, or about two percent of the sample mean of about 23 years (column (2)). Next, we add more controls and compare those who experienced occupational declines to observationally similar workers in similar occupations and industries. These specifications suggests that the losses from occupational decline averaged about two months (0.2 years) of employment, or about one percent of the sample mean (columns (3)-(6)).

Table 3.4 about here

30. Earnings data for the population of Swedish workers are not available at annual frequency prior to 1985. We obtain prior earnings from the population censuses, which were carried out every five years until 1990.

Panel B of Table 3.4 reports results from using cumulative earnings 1986-2013 as the outcome. Column (1) shows that working in a declining occupation was associated with 350,000 Swedish Krona (SEK) lower cumulative earnings, or about 5% of the sample mean.[31] When including individual controls, the estimated loss is similar, though the confidence interval is much tighter (column (2)). Further adding occupational controls cuts the loss to less than two percent of the sample mean.

In panel C we examine earnings losses from occupational decline using an alternative earnings measure: cumulative earnings divided by predicted initial earnings (see Section 3.3.2 for details on the construction of this variable). Depending on the controls included, the estimated losses in cumulative earnings range from around 100% to 220% of initial annual earnings, or from 2.5%-5.7% in terms of the sample mean, quite similar to the results in panel B.[32]

In Figure 3.2, we present a dynamic counterpart to the results reported in panel B, columns (2) and (6) of Table 3.4. Here we use as outcomes each year's earnings and cumulative earnings from 1986 up to the year indicated on the horizontal axis of each chart. The top right panel of Figure 3.2 is suggestive of a smooth process of occupational decline, with earnings losses building up gradually. However, the top left panel reveals that losses in annual earnings suffered by workers in declining occupations were larger during the 1990s and late-2000s recessions.[33] The picture is similar when we divide the coefficients by the mean of cumulative earnings at each horizon (bottom panels). As before, the losses are smaller when we include occupation and industry controls.

Figure 3.2 about here

Next, we investigate occupational mobility. Table 3.5 reports estimates of equation (3.4) with indicators for working in 2013 in the same occupation as in 1985 (or in a similar ones) as outcomes. As we do not want to condition the sample on being employed in 2013 (which is also an outcome), we have that 'not remaining' in the same occupation could reflect either occupational switching or non-employment, a point to which we return below.

Table 3.5 about here

Column (1) in panel A of Table 3.5 shows that the probability of remaining in the same 3-digit occupation was around 14 percentage points lower in declining occupations, compared to a mean of 29% in our sample. In other words, by 2013 a little over 70% of all workers had left their 1985 occupations (or left employment altogether), and the probability of staying in the same occupation was roughly halved for those starting in declining occupations. When we compare observationally similar workers, occupational decline appears to reduce the probability of remaining in the 1985 occupation by 11 percentage

31. We inflate all SEK figures to 2014 levels. Average annual earnings of Swedish workers, conditional on being employed in November and earning at least the base amount during the year, were SEK190,200 in 1985 and SEK330,800 in 2013, in terms of 2014SEK. We do not express these amounts in USD due to exchange rate fluctuations. For instance, SEK1,000 were worth about USD150 in January 2014, but about USD130 in December 2014, and about USD110 in October 2018.

32. Below we discuss results using alternative functional forms for cumulative earnings.

33. These results are related to Jaimovich and Siu (2020), who find that recessions hit routine employment particularly hard.

points, and when further restricting the comparison to similar occupations and industries, the estimate falls to 4.5 percentage points. Panels (B) and (C) of Table 3.5 show similar, albeit somewhat smaller, coefficients when we look at the probability of remaining in more broadly defined (2-digit or 1-digit) occupations. It is noteworthy that even when we consider 1-digit occupations, only about 40% of the sample remained in the same broadly defined occupation over the 28-year period that we study.[34]

Having presented our main results on career employment, career earnings, and occupational mobility, we now return to the issues of interpretation alluded to in Section 3.4.1 and the question of magnitudes discussed in Section 3.4.2. Our results from including individual, occupation, and industry controls (column (6) in Tables 3.4 and 3.5) plausibly provide us with conservative estimates of the losses from occupational decline—about 2% of mean cumulative earnings over 28 years—given the balance of pre-determined characteristics conditional on these controls, as well as the fact that no earnings losses appear in the first 5-10 years (Figure 3.2). As argued in Section 3.5.1, the specification only controlling for individual characteristics (column (2) in Tables 3.4 and 3.5) may slightly overstate the losses from occupational decline—5% of mean cumulative earnings—as it leaves minor differences in some of the pre-determined characteristics. In addition, Figure 3.2 shows earnings losses based on this specification already in the years immediately after 1985. On the other hand, we have also seen that the extent of occupational decline is much reduced when including occupation and industry controls (columns (2) and (6) in Table 3.1). This brings us to the discussion of magnitudes.

We calculate the elasticities of our outcome variables with respect to occupational employment growth in Sweden by taking the reduced form estimates reported in Tables 3.4 and 3.5 (expressed in percent of the outcome mean) and dividing them by the difference in occupational employment growth between declining and non-declining occupations reported in Table 3.1. This yields elasticities of 0.04-0.05 for cumulative employment, 0.08-0.13 for cumulative earnings, and 0.71-0.95 for remaining in the initial occupation (the ranges again refer to columns (2) and (6) in the relevant tables). Recall from the discussion in Section 3.4.2 that these numbers are likely upper bounds on the true elasticities. Nonetheless, they support the interpretation that on average, occupational decline results in a modest loss in career employment, a somewhat larger but still modest loss in career earnings, and in contrast, a strong occupational mobility response. These estimates potentially mask substantial heterogeneity, however, which we investigate in Section 3.5.3.

Robustness of main results

Our first set of robustness checks relates to the choice of functional form of occupational decline. The declining indicator is based on a 25% cutoff, conservatively identifying occupations whose (US) employment fell substantially since the mid 1980s. We also explore a range of alternative cutoffs and find that higher cutoffs (in the sense of isolating larger employment declines) usually result in larger estimated losses and mobility responses. In addition, our results are very similar when we exclude occupations that grew rapidly from the control group (see Table 3.10 for both sets of results). While we

34. For related discussions of the importance of switching occupations in the presence of technological change, see Cortes (2016) and Caselli and Manning (2019).

focus on a binary definition of occupational decline as motivated in Section 3.4.2, we also explore the relationships between our key outcomes of interest and the full variation in US and Swedish employment growth. The (residualised) associations of cumulative earnings and occupational mobility with occupational employment growth are mostly flat apart from a drop in occupations that declined substantially (see Figures 3.5 and 3.6, and corresponding regression results in Table 3.11).

A second set of robustness checks adds fixed effects for the firms that workers worked for in 1985. This specification lets us compare two workers who started out in declining and non-declining occupations, respectively, but worked at the same firm. Thus, we address the concern that declining occupations may be systematically different in terms of the quality of firms, which may differentially affect workers' future careers. Reassuringly, as Table 3.12 shows, the findings in Tables 3.4 and 3.5 are generally robust to adding fixed effects for workers' starting firms. The effects on employment losses and occupational stability are a little smaller in magnitude, while those on earnings are a little larger. This suggests that, conditional on our other controls, differential sorting into firms does not drive our estimates.

Finally, our conclusions about earnings losses are robust to using different functional forms of earnings, as we discuss further in the next sub-section.

3.5.3 Heterogeneity by within-occupation earnings rank

We now examine how employment and earnings losses from occupational decline varied by initial within-occupation earnings rank. We estimate equation (3.4) allowing the coefficient on the declining indicator to vary by earnings rank, and report the results in Table 3.6. Panel A shows that lower ranked workers suffered larger employment and earnings losses than average as a result of occupational decline (columns (1)-(6)): the coefficients on the interaction of the declining indicator with earnings rank are positive and precisely estimated. Moreover, these estimates barely change when we add occupation and industry controls over individual-level controls, though the main coefficients on the declining dummy—giving the employment and earnings loss for the median worker—are affected by the inclusion of additional controls. The magnitudes implied by the interaction coefficients are meaningful and imply, for instance, that compared to the 25th-percentile, the 75th-percentile worker suffered a 5% lower employment loss and a 6.5% lower earnings loss (both in terms of the overall mean).

Table 3.6 about here

This pattern is robust to alternative specifications that replace the linear rank measures with dummies for the bottom and top terciles. This specification also allows us to characterise losses for low-ranked workers directly. Panel B of Table 3.6 shows that workers at the bottom tercile of their starting occupations' earnings distributions suffered employment losses of 1.2-1.4 years (5.5%-6.5% of mean employment in the bottom tercile) and earnings losses of around 8%-11% of bottom-tercile mean earnings. Indeed, the estimates of mean losses reported in the previous sub-section mask more substantial losses for low earners (within an occupation). Our findings about earnings losses are robust to a number of alternative ways of measuring career earnings, as shown in Table 3.13.[35]

The pattern for the probability of remaining in the initial occupation appears to be non-monotonic: among the workers in declining occupations, both bottom-tercile and top-tercile workers were less likely to remain in their starting occupations (panel B, columns (7)-(8)). These interaction coefficients are larger than ten percent of the overall mean (although in the case of the top tercile, not precisely estimated). This hump-shaped pattern of staying probabilities (U-shaped in exiting probabilities) is intriguing from a theoretical point of view, as we discuss below.

One potential challenge in interpreting the results of Table 3.6 is that those with low earnings in their occupation may have differed from others along some observable dimensions, such as gender, age, or geography. To mitigate this concern, we re-estimate the regressions using workers' within-occupation rank in residualised earnings, where the residuals come from a regression of earnings on female, cohort, and county-of-residence dummies. As Table 3.14 shows, in terms of employment and earnings losses the results are qualitatively unchanged, and the magnitude of the interaction coefficients is only slightly reduced. However, using the residual-based rank measure, there is less support for the conclusion that bottom-ranked workers were less likely to remain in the initial occupation. Table 3.15 reports interactions using overall earnings rank instead of within-occupation earnings rank. Large losses from occupational decline for those in the lowest tercile again stand out as a consistently robust finding.[36]

We consider three further dimensions of heterogeneity. First, we examine earnings losses separately for those who remained in their initial occupation and those who did not. This purely descriptive exercise is motivated by the prediction of our baseline model in Section 3.2 that leavers should have lower losses than stayers. We estimate equation (3.4) with cumulative earnings as the outcome variable, and add on the right-hand side a dummy for having remained in the initial occupation, as well as its interaction with the declining dummy. Panel A of Table 3.17 shows that among all workers, those who remained in their initial occupation had higher cumulative earnings, though in panel B we restrict the sample to those who were employed in 2013, and the finding reverses.[37] Importantly, in neither case is there evidence that those who remained in declining occupations did significantly worse than those who left a declining occupation. The same result holds when we focus on the bottom third (in terms of within-occupation earnings), see panel C. We discuss the interpretation of these results in light of the model in Section 3.5.6 below.

Second, while we argue that our approach delivers an upper bound on the magnitude of average losses due to occupational decline, we also explore heterogeneity related to workers in non-declining occupations. Workers who leave declining occupations may flock to similar non-declining occupations, depressing the wage in these 'control' occupations. And even in the absence of such general equilibrium

35. We consider discounted cumulative earnings, applying a 5% discount rate; discounted cumulative earnings normalised by initial earnings; the percentile rank in cumulative earnings; the log of cumulative earnings; and the percentage change in earnings 1985-2013. As expected, the estimated losses in terms of discounted cumulative earnings are somewhat smaller at 1.5%-4.5% of the overall mean, depending on controls, as more weight is put on earlier years in the career.

36. Table 3.16 uses residualised overall earnings rank, again with similar results. Note that in the theoretical model, within-occupation and overall earnings rank and their residualised counterparts are all identical.

37. Workers classified as having remained are employed in 2013 by construction, whereas those classified as not having remained might not have been employed in 2013 and thus have zero earnings in that year, and possibly in preceding years also.

effects, employing a rich set of controls may cause us to put more weight on groups of comparable occupations where there are roughly as many declining as non-declining sub-occupations. In such cases, workers may have many substitute occupations to choose from. To explore such issues, we run what we refer to as 'doughnut' specifications, namely the same regressions as those we report in Tables 3.4 and 3.5 but excluding 3-digit (SSYK96) occupations in which some but not all 5-digit occupations are declining. We indeed estimate slightly larger earnings losses than in our baseline specifications, ranging from 3%-6% of mean earnings (see Tables 3.18 and 3.19).

Third and finally, we explore heterogeneity by gender. As Table 3.20 shows, occupational decline results in larger losses of employment and occupational stability for women, while men suffer larger earnings losses. Larger employment and earnings losses for those in the bottom tercile are concentrated among men, although once we use residualised earnings rank (Table 3.21) the losses of women in the bottom tercile are roughly as large as men's.

3.5.4 Unemployment, retraining, early retirement, and geographic stability

A natural question at this stage is to what extent the loss in years of employment due to occupational decline is accounted for by increased unemployment and retraining; as discussed above, data on these last two outcomes are available for the final 22 years of our study. Table 3.7 reports estimates using the main specifications from Tables 3.4 and 3.6 but this time using cumulative days of unemployment (panel A) and state-sponsored retraining (panel B) as outcome variables. Columns (1)-(4) of Panel A show that workers who started out in later declining occupations were only very slightly more likely to ever be unemployed, and columns (5)-(8) suggest that these workers accumulated 20-50 more unemployment days, though the estimates with more controls are imprecise. However, we again find substantial heterogeneity, with bottom-tercile workers in declining occupations spending 63 days more in unemployment, a substantial 20% of the mean.

Table 3.7 about here

Columns (1)-(4) of panel B suggest that occupational decline increased the risk of ever enrolling in state-sponsored retraining by 9%-27%. The estimates for cumulative days spent retraining are similarly substantial, at least in relative terms (columns (5)-(8)). Our most conservative specification including all controls suggest that the median worker spent six more days in retraining, which amounts to 21% of the mean (ten days and 29% for the bottom-tercile worker).

Our estimates for unemployment and retraining can only explain part of the estimated employment losses. For bottom-tercile workers, we conservatively estimate an employment loss of 1.16 years.[38] Of these, unemployment and retraining account for only 22%.[39] The remaining employment loss may be

38. From panel B, column (2) in Table 3.6 we obtain $-0.03 - 1.13 = -1.16$. To complete the calculation, we divide the unemployment and retraining coefficients by 365 to get years, multiply them by $28/22$ to account for the fact that these variables are only available during 1992-2013, sum them, and divide by 1.16.

39. Of the mean employment loss, unemployment and retraining explain about a third and a tenth of the time respectively.

accounted for by job search that is not covered by unemployment benefits; private retraining; or time spent outside the labour force. Unfortunately, we lack the data to investigate this further.

There is however a group of workers for whom we are able to investigate the relationship between occupational decline and exit from the labour force, namely, older workers. Recall that workers in our baseline sample reached a maximum age of 64 in 2013. We now examine employment, earnings, and retirement for two groups of older workers, most of whom reached the usual retirement age of 65 well before the end of our sample period.

Panel A of Table 3.22 considers workers who were aged 37-48 in 1985. The employment losses among this group are a little larger than for our baseline sample: about 8 months (4 months) of a year of employment in the specification with individual (all) controls, or just under 4% (2%) of the group mean. About half of these employment losses are accounted for by a slightly younger age of retirement for those in declining occupations. The estimated earnings losses from occupational decline—about 6% (1.5%) with individual (all) controls—are similar to those of the baseline group. Finally, for this group we also find positive and significant interactions of the declining dummy with initial occupational earnings rank, suggesting once more that those who earned least within their occupation to begin with lost more years of employment from occupational decline.

Panel B of Table 3.22 suggests that for an even older group, those aged 49-60 in 1985, the occupational decline that we measure had more modest costs compared to the baseline group. This likely reflects the fact that we are measuring occupational decline over a longer period, and that these older workers had little exposure to the decline.[40]

Finally, we study whether occupational decline results in geographic mobility across municipalities, commuting zones, and counties. The results in Tables 3.23 and 3.24 suggest that occupational decline does not, on average, reduce geographic stability. Workers in the bottom tercile, however, are a little more likely to leave their location when their occupation declines.

3.5.5 Technology-related occupational decline

Consistent with much of the literature (Goos, Manning, and Salomons 2014) we expect technological change to be a key driver of occupational decline, and especially occupational decline that is common to the US and Sweden. Nevertheless, there could be other drivers, including changes on the supply side (changes in demographics, trade shocks, or changes in government policy) and in consumer demand. Bearing this in mind, we now focus on occupations that are likely to have declined due to the introduction of labour-replacing technology, based on information from the OOH, as described in Section 3.3.2.[41]

We find that workers' exposure to declines in Swedish occupational employment is of very similar magnitude regardless of whether we consider all occupations classified as declining, or only the ones

40. We verify that for the groups of middle-aged and older workers, our declining indicator does not predict differences in prior earnings (1975 and 1980) conditional on controls, see Figure 3.4.

41. Some of what we classify as technology-related decline may still be influenced by other factors, and we cannot rule out that technology played a role in the remaining declining occupations.

we linked to technology (Table 3.26, panel A, and Table 3.27).[42] Moreover, technology-related occupational declines are not significantly different from other occupational declines in their implications for years of employment, cumulative earnings, and the probability of remaining in the initial occupation. One way to see this is by adding an indicator for technology-linked decline to equation (3.4). We find that the coefficients on this indicator are statistically indistinguishable from zero (columns (1) and (2) in panels B-D of Table 3.26). Alternatively, considering technology-related declines on their own, we see very similar point estimates, both for the main effect and for the interaction with earnings rank, as for the full set of declines (columns (3)-(5)).

Our second approach to investigating the consequences of technology-driven occupational decline relies on the presence of relevant labour-replacing technologies, classified using our algorithmic or manual approaches (as described in Section 3.3.2). Panel A of Table 3.28 shows that both measures of technological replacement are correlated with occupational decline, although the estimates for the algorithmic measure is a little larger and (once we include all controls) more precisely estimated.[43] As Panel B shows, both measures of technological replacement are also correlated with large employment declines at the coarse 3-digit level, although only the algorithmic measure survives the inclusion of the full set of controls. Panel C shows that at the individual level, both measures result in fairly moderate employment losses, with point estimates that are all below half a year of employment—similar to our main measure of occupational decline. Finally, Panel D shows that the cumulative earnings losses from technological replacement are also quite similar to our main estimates: around 1%-3% (imprecisely estimated) for the algorithmic measure, and around 5%-7% (precisely estimated) for the manual measure.

3.5.6 Interpreting our findings through the lens of the theoretical model

In our empirical analysis, we confirm that occupational decline was associated with earnings losses and higher occupational exit rates. This is consistent with our model's assumption that occupational decline was largely driven by changes in demand.

The version of the model which best fits our empirical findings is the one with both differential occupational switching costs and displacement. In this case, those with lower initial within-occupation earnings rank suffer larger earnings losses as a result of occupational decline; switchers' earnings losses may be larger than those of stayers (as we find); and displacement may lead to switching probabilities that are U-shaped in initial earnings, whereby low-earning workers switch if displaced, while high-earning workers switch voluntarily.

Our empirical analysis also sheds light on the nature of the occupational switching costs in the model. In practice we find that roughly a third of the employment years lost can be accounted for by increased unemployment, and almost ten percent are due to retraining. The stronger responses to occupational decline of unemployment and retraining among lower-ranked workers further supports our

42. Workers starting out in 1985 in subsequently declining occupations, where we were able to identify a link to technology, were statistically indistinguishable from those in the remaining declining occupations, as seen in Table 3.25.

43. Both measures are (conditionally) balanced on the same set of characteristics that we report in Table 3.3—these estimates are available on request.

interpretation of heterogeneous switching costs.

Finally, our model suggests that the effects of an adverse occupational demand shock may differ depending on whether the shock was anticipated. Controlling for projected employment growth—thus isolating unanticipated declines—generally leads us to estimate smaller earnings losses and mobility responses. The former is consistent with our model, but the latter is not. A possible explanation may be that conditional on predicted occupational employment growth, our declining indicator isolates a lower level of exposure to actual Swedish employment declines than in the unconditional regression, or the one only conditioning on individual characteristics.[44]

3.6 Conclusion

In this paper, we study the long-run employment and earnings losses that workers suffer when demand for their occupations declines. We begin by measuring anticipated and actual occupational declines in the US, which we map into panel micro data on Swedish workers. We find that even after controlling for key predictors of occupational decline, employment changes in declining Swedish occupations were around 20-40 log points lower than in non-declining occupations.

Despite this large fall in employment, we find that over 28 years, those who in 1985 worked in declining occupations experienced earnings (employment) losses that were around 2%-5%(1%-2%) of mean cumulative earnings (employment), compared to those who initially worked in non-declining occupations. We characterise conditions under which these figures are a plausible upper bound on the magnitude of the losses due to occupational decline. The earnings losses are on the higher end of the above-mentioned range when we control only for individual covariates, and lower when we also control for anticipated occupational changes and industry and occupation characteristics. Around a third of the cumulative employment losses are accounted for by increased unemployment, and a further tenth by increased time spent in government retraining.

We find that workers in the bottom tercile of their occupations' earnings distributions suffered the largest losses (around 8%-11%). Workers in the bottom tercile also lost more years of employment and spent more time in unemployment and retraining. We find that those in declining occupations were significantly more likely to leave their starting occupations. The propensity to exit declining occupations was U-shaped in initial occupational earnings rank, with those at the bottom (and to a lesser extent at the top) more likely to leave their starting occupations.

We show that our findings are consistent with a Roy model with negative occupational demand shocks, where workers may suffer displacement, and where finding reemployment takes time. In the model, those at the bottom of a declining occupation also have low earnings capacity in other occupations, and therefore find it harder to find reemployment—whether in their own occupations or in other occupations. Hence they lose most from occupational decline. The model also rationalises the U-shaped exit pattern that we describe above: those at the bottom of their occupations' earnings distributions are more likely to leave their occupations when they are displaced, while those at the top are more likely to

44. However, exposure declines by less than the mobility response, in relative terms. See columns (2) and (4) in Tables 3.1 and 3.4.

leave to avoid negative demand shocks.

Our findings suggest that the mean losses of occupational decline are lower than the losses suffered by displaced workers that have been reported in prior literature. This is likely because occupational decline is typically gradual, and can be partly managed through retirements, reduced entry into declining occupations, and increased job-to-job exits to other occupations. Indeed, we document a response in terms of occupational mobility which, in terms of elasticities, is much larger than our estimated earnings losses. Gradual occupational decline may also impose fewer negative spillovers on local economies compared to large, sudden shocks, such as plant closures.

At the same time, future occupational decline could still have substantial adverse consequences for workers' outcomes, for the following three reasons. First, our paper studies occupational decline that—while unanticipated early in workers' careers—was nevertheless fairly gradual. But if, for example, machine learning improves rapidly, occupational replacement may happen faster, and may be accompanied by an overall worsening of employment opportunities (Bostrom 2014). Second, the occupational decline that we study largely spared the most skilled occupations, but this may change with new technologies. Many professionals made sizeable investments in skills that are particularly useful in their occupations, and some may also benefit from economic rents. It is possible that for these workers the earnings losses from future occupational decline may be higher than those we estimate. Finally, and perhaps most importantly, our findings show that low-earning individuals are already suffering considerable (pre-tax) earnings losses, even in Sweden, where institutions are geared towards mitigating those losses and facilitating occupational transitions. Helping these workers stay productive when they face occupational decline remains an important challenge for governments.

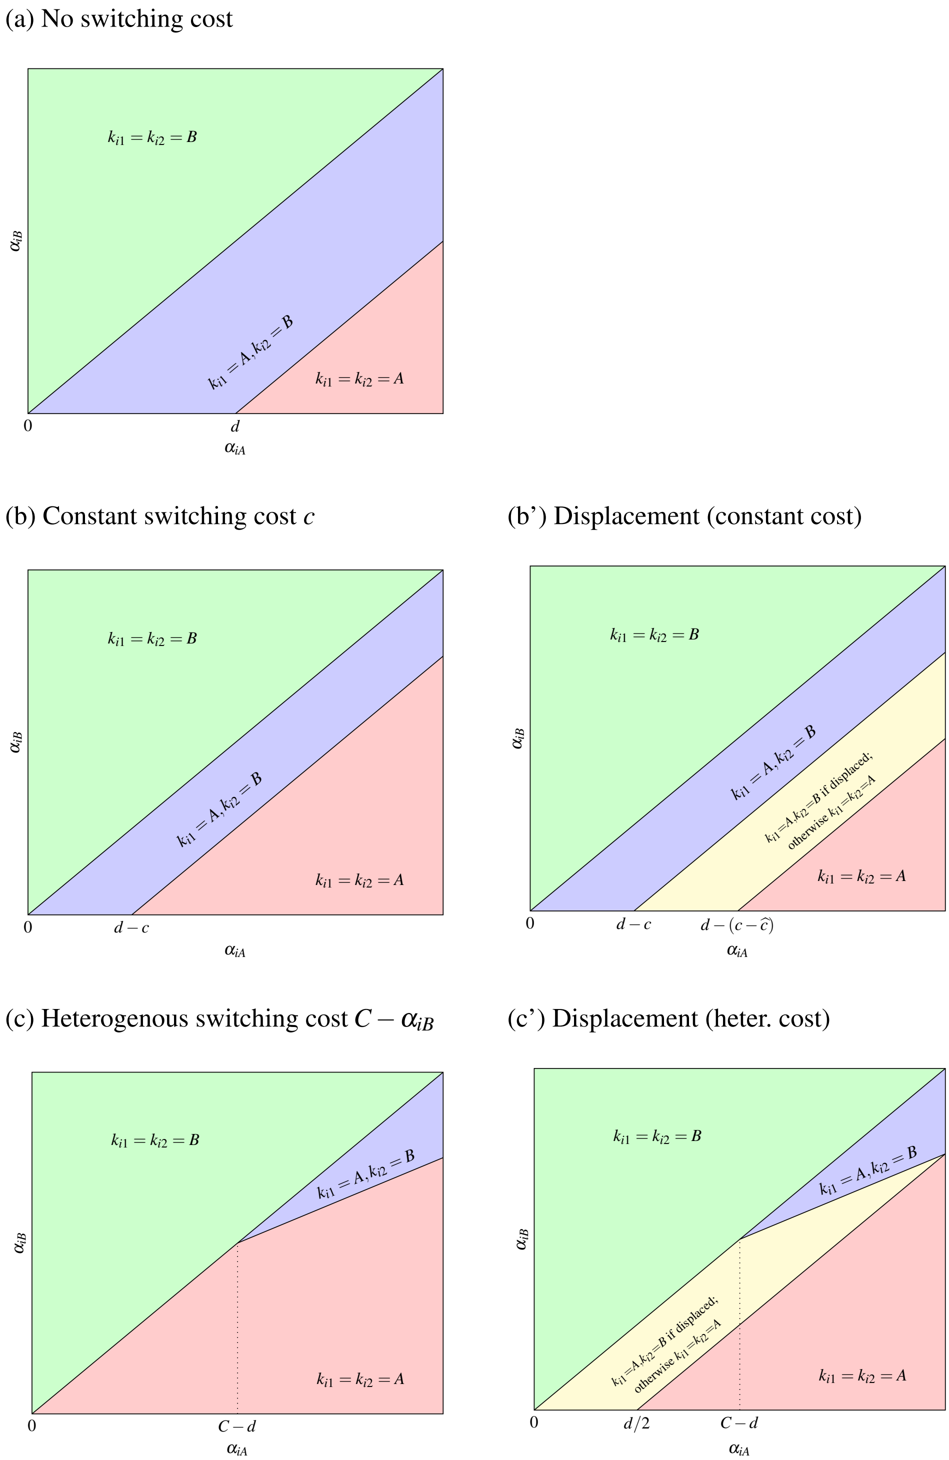

Notes: k_{it} denotes the occupation chosen by worker i in period t. α_{ik} denotes log productivity of worker i in occupation k. d is the amount by which the relative occupational log price declines from period 1 to period 2. The parameter values chosen are $(\overline{\alpha}, d, c, \widehat{c}, C) = (1, 0.5, 0.25, 0.25, 1)$.

Figure 3.1: Sorting in a two-period Roy model

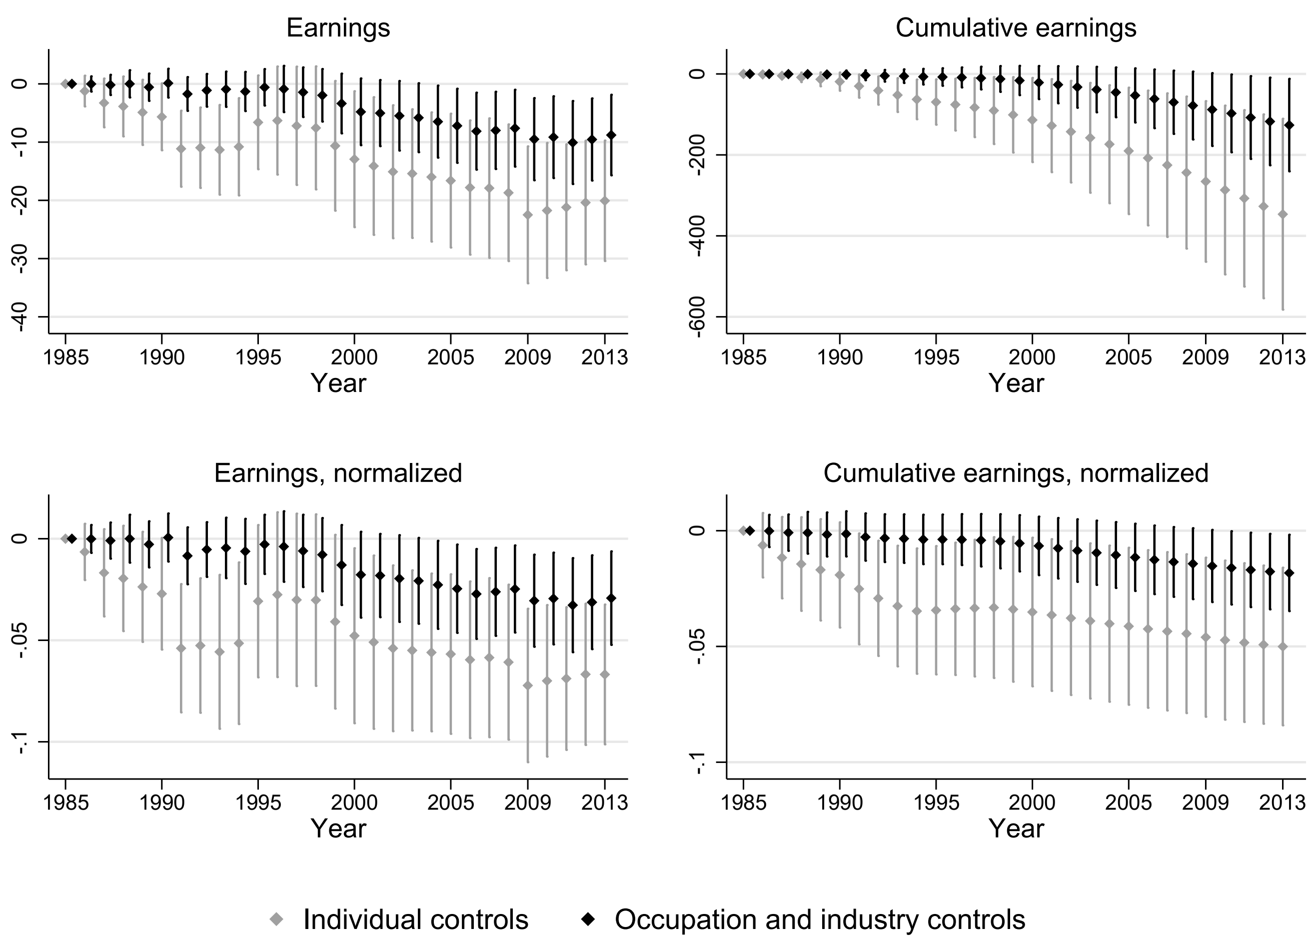

Notes: Diamonds mark the coefficients on the declining indicator from regressions of annual earnings or cumulative earnings on the indicator, including the same controls as in columns (2) ('Individual controls') and (6) ('Occupation and industry controls') of Table 3.4, separately for each year 1986-2013. Bars indicate 95-percent confidence intervals.

Figure 3.2: Differences in earnings and cumulative earnings by exposure to occupational decline, over time

Bibliography for Chapter 3

Acemoglu, Daron, and Pascual Restrepo. 2019. “Automation and New Tasks: How Technology Displaces and Reinstates Labor.” *Journal of Economic Perspectives* 33 (2): 3–30.

———. 2022. “Tasks, Automation, and the Rise in U.S. Wage Inequality.” *Econometrica* 90 (5): 1973–2016.

Adermon, Adrian, and Magnus Gustavsson. 2015. “Job Polarization and Task-Biased Technological Change: Evidence from Sweden, 1975-2005.” *Scandinavian Journal of Economics* 117 (3): 878–917.

Andersson, Josefine. 2017. “Insurances against job loss and disability: Private and public interventions and their effects on job search and labor supply.” PhD diss., Uppsala University.

Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. 2017. “Revisiting the risk of automation.” *Economics Letters* 159:157–160.

Autor, David, and David Dorn. 2009. “This Job Is ”Getting Old”: Measuring Changes in Job Opportunities Using Occupational Age Structure.” *American Economic Review* 99 (2): 45–51.

Autor, David H. 2015. “Why Are There Still So Many Jobs? The History and Future of Workplace Automation.” *Journal of Economic Perspectives* 29 (3): 3–30.

Autor, David H., David Dorn, Gordon H. Hanson, and Jae Song. 2014. “Trade Adjustment: Worker-Level Evidence.” *The Quarterly Journal of Economics* 129 (4): 1799–1860.

Bachmann, Ronald, Merve Cim, and Colin Green. 2019. “Long-Run Patterns of Labour Market Polarization: Evidence from German Micro Data.” *British Journal of Industrial Relations* 57 (2): 350–376.

Bana, Sarah H. 2021. “Identifying Vulnerable Displaced Workers:The Role of State-Level Occupation Conditions.” Working paper.

Battisti, Michele, Christian Dustmann, and Uta Schönberg. 2017. “Technological and Organizational Change and the Careers of Workers.” Working paper.

Bessen, James. 2016. *How computer automation affects occupations: Technology, jobs, and skills.* Law and Economics Research Paper 15-49. Boston University School of Law.

Bo, Ernesto Dal, Frederico Finan, Olle Folke, Torsten Persson, and Johanna Rickne. 2022. “Economic and Social Outsiders but Political Insiders: Sweden’s Populist Radical Right.” *The Review of Economic Studies.*

Bostrom, Nick. 2014. *Superintelligence: Paths, Dangers, Strategies.* Oxford University Press.

Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies.* W. W. Norton & Company.

Bureau of Labor Statistics. 1986. *Occupational Outlook Handbook, 1986-87 Edition: Bulletin of the United States Bureau of Labor Statistics, No. 2250.* US Department of Labor.

———. 2018. *Occupational Outlook Handbook, 2018-2019 Edition.* Bernan Press.

Campbell, Peter. 2018. "Trucks headed for a driverless future." *Financial Times* January 31.

Caprettini, Bruno, and Hans-Joachim Voth. 2017. *Rage Against the Machines: Labour-Saving Technology and Unrest in England, 1830-32.* CEPR Discussion Papers DP11800. Centre for Economic Policy Research, January.

Caselli, Francesco, and Alan Manning. 2019. "Robot Arithmetic: New Technology and Wages." *American Economic Review: Insights* 1 (1): 1–12.

Cortes, Guido Matias. 2016. "Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data." *Journal of Labor Economics* 34 (1): 63–105.

Cortes, Guido Matias, and Giovanni Gallipoli. 2017. "The costs of occupational mobility: An aggregate analysis." *Journal of the European Economic Association* 16 (2): 275–315.

Dauth, Wolfgang, Sebastian Findeisen, Jens Suedekum, and Nicole Woessner. 2021. "The Adjustment of Labor Markets to Robots." *Journal of the European Economic Association* 19 (6): 3104–3153.

Davis, Steven, and Till Von Wachter. 2011. "Recessions and the Costs of Job Loss." *Brookings Papers on Economic Activity* 42 (2 (Fall)): 1–72.

Edin, Per-Anders, Tiernan Evans, Georg Graetz, Sofia Hernnäs, and Guy Michaels. 2019. *Individual Consequences of Occupational Decline.* CEPR Discussion Papers 13808. C.E.P.R. Discussion Papers, June.

Edin, Per-Anders, and Bertil Holmlund. 1995. "The Swedish Wage Structure: The Rise and Fall of Solidarity Wage Policy?" In *Differences and Changes in Wage Structures,* 307–344. NBER Chapters. National Bureau of Economic Research, Inc, October.

Eliason, Marcus, and Donald Storrie. 2006. "Lasting or Latent Scars? Swedish Evidence on the Long-Term Effects of Job Displacement." *Journal of Labor Economics* 24 (4): 831–856.

Frey, Carl Benedikt, and Michael A. Osborne. 2017. "The future of employment: How susceptible are jobs to computerisation?" *Technological Forecasting and Social Change* 114:254–280.

Galaasen, Sigurd Mølster, and Andreas Ravndal Kostøl. 2018. "Mismatch and the Consequences of Job Loss." Working paper.

Gathmann, Christina, Ines Helm, and Uta Schönberg. 2018. "Spillover effects in local labor markets: Evidence from mass layoffs." *Journal of the European Economic Association* 18 (1): 427–468.

Gathmann, Christina, and Uta Schönberg. 2010. "How general is human capital? A task-based approach." *Journal of Labor Economics* 28 (1): 1–49.

Gola, Paweł. 2021. "Supply and demand in a two-sector matching model." *Journal of Political Economy* 129 (3): 940–978.

Goldin, Claudia. 2000. "Labor Markets in the Twentieth Century." In *The Cambridge History of the United States,* edited by Robert E. Gallman Stanley L. Engerman, 3:549–624. Cambridge University Press.

Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review* 104 (8): 2509–26.

Gottfries, Nils. 2018. "The labor market in Sweden since the 1990s." *IZA World of Labor,* 411–411.

Graetz, Georg. 2020. *Technological Change and the Swedish Labor Market.* Working paper 2020:19. IFAU.

Humlum, Anders. 2021. *Robot Adoption and Labor Market Dynamics.* Working paper. University of Chicago.

Jacobson, Louis S., Robert J. LaLonde, and Daniel G. Sullivan. 1993. "Earnings Losses of Displaced Workers." *The American Economic Review* 83 (4): 685–709.

Jaimovich, Nir, and Henry E. Siu. 2020. "Job Polarization and Jobless Recoveries." *The Review of Economics and Statistics* 102 (1): 129–147.

Kambourov, Gueorgui, and Iourii Manovskii. 2009. "Occupational specificity of human capital." *International Economic Review* 50 (1): 63–115.

Lindbeck, Assar. 1997. "The Swedish Experiment." *Journal of Economic Literature* 35 (3): 1273–1319.

Lindqvist, Erik, and Roine Vestman. 2011. "The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment." *American Economic Journal: Applied Economics* 3 (1): 101–128.

Marx, Karl. 1867. *Das Kapital.* Verlag von Otto Meisner.

Neal, Derek. 1995. "Industry-Specific Human Capital: Evidence from Displaced Workers." *Journal of labor Economics* 13 (4): 653–677.

OECD. 2015. *Back to Work: Sweden.*

Parent, Daniel. 2000. “Industry-Specific Capital and the Wage Profile: Evidence from the National Longitudinal Survey of Youth and the Panel Study of Income Dynamics.” *Journal of Labor Economics* 18 (2): 306–323.

Pavan, Ronni. 2011. “Career Choice and Wage Growth.” *Journal of Labor Economics* 29 (3): 549–587.

Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory, 2nd Edition.* The MIT Press.

Poletaev, Maxim, and Chris Robinson. 2008. “Human capital specificity: evidence from the Dictionary of Occupational Titles and Displaced Worker Surveys, 1984–2000.” *Journal of Labor Economics* 26 (3): 387–420.

Roy, A. D. 1951. “Some Thoughts on the Distribution of Earnings.” *Oxford Economic Papers* 3 (2): 135–146.

Scheiber, Noam. 2018. “High-Skilled White-Collar Work? Machines Can Do That, Too.” *New York Times* July 07.

Schmillen, Achim. 2018. “Vocational education, occupational choice and unemployment over the professional career.” *Empirical Economics.*

Skans, Oskar Nordström, Per-Anders Edin, and Bertil Holmlund. 2009. “Wage dispersion between and within plants: Sweden 1985-2000.” In *The structure of wages: An international comparison,* 217–260. University of Chicago Press.

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

Statistics Sweden. 1990. *Trender och prognoser: med sikte på 2015* [in Swedish]. Stockholm: Statistics Sweden.

Sullivan, Paul. 2010. “Empirical Evidence on Occupation and Industry Specific Human Capital.” *Labour Economics* 17 (3): 567–580.

Susskind, Richard E, and Daniel Susskind. 2015. *The Future of the Professions: How Technology Will Transform the Work of Human Experts.* Oxford University Press.

Veneri, Carolyn M. 1997. “Evaluating the 1995 occupational employment projections.” *Monthly Labor Review* 120.

Vikström, Johan, and G van den Berg. 2017. *Långsiktiga effekter av arbetsmarknadsutbildning.* Technical report. IFAU Rapport 2017:17.

Table 3.1: Quantifying workers' exposure to occupational decline

	(1)	(2)	(3)	(4)	(5)	(6)
A. Workers aged 16-64 in 1985 (3,061,051 observations)						
Declining	-0.49	-0.44	-0.43	-0.31	-0.28	-0.22
	(0.12)	(0.11)	(0.11)	(0.10)	(0.11)	(0.10)
B. Workers aged 25-36 in 1985 (877,324 observations)						
Declining	-0.47	-0.40	-0.39	-0.28	-0.27	-0.22
	(0.11)	(0.11)	(0.11)	(0.10)	(0.12)	(0.10)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of occupational log employment changes on a dummy for working in a declining occupation are shown. Regressions are run on individual-level data. However, the dependent variable is the difference in aggregate log employment in Swedish 3-digit occupations between 2013 and 1985, matched to each workers' 1985 5-digit occupation using a cross-walk. A Swedish 5-digit occupation is classified as 'Declining' if there are employment losses of more than 25 percent between 1986-2016 in the corresponding US occupation(s). In the regressions reported here, the 'Declining' variable indicates that an individual worked in such an occupation in 1985. Demographic controls include female, cohort, county, and education dummies. Occupation-level life-cycle profiles are cumulative earnings calculated for each individual based on their 1985 occupation. Predictors of growth include 1985 employment shares, 1960-85 occupational employment growth, and the predicted growth index. Occupation and industry dummies are at the 1-digit and 2-digit levels, respectively. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.2: Baseline characteristics of workers in subsequently declining occupations

	(1) Female	(2) age	(3) Compulsory school	(4) High school	(5) College	(6) Earnings	(7) Manufacturing
A. Workers aged 16-64							
Intercept	0.52	39.5	0.33	0.56	0.11	191.3	0.25
	(0.078)	(0.41)	(0.030)	(0.033)	(0.027)	(10.8)	(0.050)
Declining	-0.25	-0.89	0.13	-0.063	-0.070	-0.23	0.38
	(0.088)	(0.63)	(0.035)	(0.034)	(0.028)	(11.0)	(0.085)
B. Workers aged 25-36							
Intercept	0.51	30.8	0.23	0.64	0.13	182.8	0.23
	(0.078)	(0.078)	(0.022)	(0.033)	(0.032)	(9.28)	(0.050)
Declining	-0.26	-0.19	0.15	-0.065	-0.082	12.0	0.38
	(0.085)	(0.091)	(0.030)	(0.034)	(0.034)	(9.40)	(0.084)

Notes: Results from OLS regressions of various baseline (1985) characteristics on a constant and an indicator for working in a declining occupation are shown (see the notes to Table 3.1 for the definition of the declining indicator). Earnings are measured in thousand Swedish crowns inflated to 2014 levels. The sample includes all individuals of the indicated ages who were employed and earned at least the base amount in 1985, and whose education, occupation, and industry are observed. The number of observations is 3,061,051 in panel A and 877,324 in panel B. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.3: Balance of pre-determined characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
A. Military test scores						
	Cognitive skills			Non-cognitive skills		
Declining	-0.24	-0.015	0.022	-0.20	-0.077	-0.022
	(0.084)	(0.035)	(0.022)	(0.062)	(0.041)	(0.021)
Individual controls		✓	✓		✓	✓
Occupation & industry controls			✓			✓
Mean of dep. var.		0.06			0.06	
Observations	272,350					
B. Mother's characteristics						
	Mother finished high school			Mother's earnings (1985)		
Declining	-0.059	-0.012	0.0033	-6.73	-2.31	0.079
	(0.020)	(0.0100)	(0.0058)	(2.41)	(1.29)	(0.84)
Individual controls		✓	✓		✓	✓
Occupation & industry controls			✓			✓
Mean of dep. var.		0.35			97.4	
Observations	609,075					
C. Father's characteristics						
	Father finished high school			Father's earnings (1985)		
Declining	-0.069	-0.0088	0.0075	-13.7	-1.38	2.26
	(0.027)	(0.012)	(0.0067)	(6.67)	(2.85)	(1.99)
Individual controls		✓	✓		✓	✓
Occupation & industry controls			✓			✓
Mean of dep. var.		0.43			174.2	
Observations	451,301					

Notes: Results from regressions of various pre-determined characteristics on a dummy for working in 1985 in a subsequently declining occupation are shown. Test scores from the military enlistment are standardised to have mean zero and unit variance within enlistment cohorts. The sample in panel A includes men born in Sweden from 1952-1959 with non-missing test scores (more than 85 percent of men in each cohort), who were employed and earned at least the base amount in 1985, and whose education, occupation, and industry are observed. The samples in panels B and C are the same as that in panel B of Table 3.2, except that individuals with missing information on mother's or father's education and income were dropped. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.4: Occupational decline and individual-level cumulative employment and earnings 1986-2013

	(1)	(2)	(3)	(4)	(5)	(6)
A. Cumulative years employed 1986-2013 (mean: 23.4)						
Declining	-0.73	-0.49	-0.49	-0.30	-0.24	-0.19
	(0.26)	(0.20)	(0.20)	(0.20)	(0.18)	(0.14)
B. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,926)						
Declining	-354	-347	-241	-117	-63	-126
	(419)	(120)	(81)	(76)	(71)	(58)
C. Cumulative real earnings divided by predicted initial earnings (mean: 38.7)						
Declining	-4.29	-2.10	-2.21	-1.52	-0.98	-1.11
	(0.91)	(0.53)	(0.54)	(0.54)	(0.41)	(0.36)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1985 in a subsequently declining occupation are shown. Demographic controls include female, cohort, county, and education dummies. Occupation-level life-cycle profiles are cumulative earnings calculated for each individual based on their 1985 occupation. Predictors of growth include 1985 employment shares, 1960-85 occupational employment growth, and the predicted growth index. Occupation and industry dummies are at the 1-digit and 2-digit levels, respectively. The number of observations is 877,324. The sample is the same as that in panel B of Table 3.2. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.5: Occupational decline and individual occupational stability

	(1)	(2)	(3)	(4)	(5)	(6)
A. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.29)						
Declining	-0.14	-0.11	-0.11	-0.065	-0.086	-0.045
	(0.043)	(0.041)	(0.042)	(0.032)	(0.035)	(0.020)
B. Probability of working in same 2-digit occupation in 2013 as in 1985 (mean: 0.35)						
Declining	-0.12	-0.088	-0.087	-0.051	-0.071	-0.037
	(0.034)	(0.034)	(0.035)	(0.030)	(0.030)	(0.019)
C. Probability of working in same 1-digit occupation in 2013 as in 1985 (mean: 0.40)						
Declining	-0.098	-0.070	-0.069	-0.039	-0.060	-0.031
	(0.030)	(0.031)	(0.032)	(0.029)	(0.027)	(0.018)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1985 in a subsequently declining occupation are shown. See the notes to Tables 3.1 and 3.4 for the definition of the declining indicator and a description of control variables, respectively. The number of observations is 553,169. The sample is the same as that in panel B of Table 3.2, except that individuals who were employed in 2013 but not sampled in the Wage Structure Statistics had to be excluded, as it is unknown whether they work in the same occupation in 2013 as in 1985. Sampling weights are applied. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.6: Heterogeneity by within-occupation earnings rank

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Linear interaction								
Declining	-0.51	-0.23	-353.5	-131.0	-2.16	-1.19	-0.11	-0.046
	(0.21)	(0.15)	(110.7)	(55.8)	(0.55)	(0.37)	(0.041)	(0.020)
Declining $\times$ rank	1.17	1.17	441.5	449.2	2.63	2.63	-0.011	-0.00090
	(0.34)	(0.30)	(142.3)	(146.8)	(0.58)	(0.57)	(0.023)	(0.017)
B. Dummy interactions								
Declining	-0.32	-0.031	-323.2	-98.0	-1.94	-0.97	-0.083	-0.022
	(0.24)	(0.18)	(123.8)	(66.7)	(0.54)	(0.41)	(0.045)	(0.021)
Declining $\times$ bottom tercile	-1.12	-1.13	-341.8	-350.1	-2.10	-2.06	-0.046	-0.040
	(0.35)	(0.33)	(106.7)	(101.5)	(0.54)	(0.51)	(0.014)	(0.013)
Declining $\times$ top tercile	0.54	0.55	232.3	235.1	1.37	1.40	-0.047	-0.029
	(0.20)	(0.16)	(135.8)	(132.1)	(0.43)	(0.48)	(0.027)	(0.018)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Mean of dep. var.	23.4		6,926		38.7		0.29	
Mean of dep. var., bottom	22.3		6,001		35.6		0.27	
Observations	877,324						553,787	

Notes: Results from regressions of the indicated outcomes on the declining indicator, within-occupation earnings rank or tercile dummies (coefficients omitted from table), and their interactions are shown. Within-occupation earnings ranks are computed in 1985 and re-scaled so as to range from -1 to 1. In panel A, the main effect on the declining indicator thus applies to the individual earning the median income within her occupation, and the coefficient on the interaction gives the inter-quartile range. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. normalised earnings are cumulative earnings divided by initial predicted earnings. The sample for columns (1)-(6) is the same as that in Table 3.4, and for columns (7)-(8) it is the same as that in Table 3.5. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.7: Occupational decline and the incidence of unemployment and retraining

	Ever				Cumulative days			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Unemployment								
Declining	0.041	0.013	0.015	0.019	52.4	17.9	20.8	20.5
	(0.021)	(0.013)	(0.012)	(0.015)	(24.8)	(14.0)	(14.0)	(18.2)
Declining × rank			-0.036				-63.8	
			(0.012)				(21.5)	
Declining × bottom tercile				0.017				42.4
				(0.012)				(18.3)
Declining × top tercile				-0.033				-43.7
				(0.012)				(17.0)
Mean of dep. var.		0.39				262		
Mean of dep. var., bottom		0.43				317		
B. Retraining								
Declining	0.035	0.012	0.013	0.015	11.4	4.73	5.04	5.81
	(0.010)	(0.0064)	(0.0063)	(0.0081)	(2.68)	(1.46)	(1.48)	(2.26)
Declining × rank			-0.027				-8.63	
			(0.0070)				(1.98)	
Declining × bottom tercile				0.014				4.38
				(0.0072)				(2.28)
Declining × top tercile				-0.022				-6.96
				(0.0064)				(2.12)
Mean of dep. var.		0.13				29		
Mean of dep. var., bottom		0.15				35		
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓	✓	✓		✓	✓	✓

Notes: Results from regressions of the indicated outcomes on the declining indicator, within-occupation earnings rank or tercile dummies (coefficients omitted from table), and their interactions are shown. Incidence of unemployment and retraining are measured during the period 1992-2013. The sample is the same as that in panel B of Table 3.2. See the notes to Table 3.6 for a description of right-hand side variables. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Appendix

Appendix 3.A Theory appendix

Here we provide a self-contained exposition of the theoretical model discussed in Section 3.2 of the paper, including all formal derivations. We consider two occupations, one of which is hit by a negative demand shock. We investigate how workers' likelihood of leaving the affected occupation, and their earnings losses, depend on their initial earnings. Starting from a standard frictionless Roy (1951) model, we successively introduce positive and potentially heterogeneous costs of switching occupation; as well as the possibility that workers are displaced from their jobs and incur a cost to find a new job even when remaining in their initial occupation. Finally, we consider how workers' sorting differs when the negative demand shock is anticipated.

Subsection 3.A.1 Setting

We consider a competitive economy with a continuum of individuals indexed by i who live for two periods $t \in \{1,2\}$ and each supplies a unit of labour inelastically each period. There are two occupations indexed by $k \in \{A,B\}$ for the workers to choose from. Workers' per-period log earnings are given by $y_{ikt} = \pi_{kt} + \alpha_{ik} - c_{ikt}$ where π_{kt} is the time-varying and stochastic (log) price of a unit of output in occupation k, α_{ik} is the time-invariant (log) amount of output that worker i produces in occupation k, and $c_{ikt} \geq 0$ is a time cost related to occupational switching, which we discuss below.[45] There are no saving opportunities and earnings are consumed immediately. We define the life-time expected utility function as $\mathbb{E}[y_{ik1} + \beta y_{ik2}]$, where $\beta > 0$ is a discount factor. In each period, workers choose the occupation that maximises their expected utility. As a normalisation, we assume that workers always choose occupation A if indifferent. Since we focus our analysis on relative wages, we define $\widetilde{\pi}_t \equiv \pi_{Bt} - \pi_{At}$ and assume for simplicity that $\widetilde{\pi}_1 = 0$.[46] Prices are determined in equilibrium by supply and demand. However, here we take them as given, and analyse the consequences of a change to prices occurring in period 2 for occupational sorting and earnings. Note that the second period may be interpreted as all periods following this change, so β could be larger than one. For simplicity, we assume that α_{iA} and α_{iB} are independent and both uniformly distributed between zero and some finite but possibly large number $\overline{\alpha}$. We explain in the following subsections that our main results are robust to alternative distributional assumptions.

45. The time cost may reflect search or retraining (or both); we assume throughout that a worker's wage equals the value of her marginal product, $e^{\pi_{kt}+\alpha_{ik}}$. We thus abstract from any job-level rents that may arise in the presence of search frictions.

46. We do not claim to identify any aggregate gains from technological change, and we do not model them here.

In period 2, there is a negative demand shock to occupation A such that $\pi_{A2} - \pi_{A1} = -d$ and $\widetilde{\pi}_2 = d, d > 0$. This may be due to labour-replacing technology becoming available, or cheaper, in occupation A. We are interested in the consequences of the shock for the earnings of workers who start out in occupation A, under various assumptions about switching costs and anticipation of the price change. Formally, let $l_i \equiv \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 1]$ be the expected earnings loss in period 2 that results from the shock, conditional on worker i starting out in occupation A, and conditional on her ability (and hence earnings rank) α_{iA}, where the occurrence of the shock is indicated by $D_A \in \{0,1\}$. Similarly, l_i^{switch} and l_i^{stay} denote the earnings losses further conditioned on leaving and staying, respectively, and p_i is the probability of switching.[47] The overall loss is given by

$$l_i = l_i^{\text{stay}} - p_i \left(l_i^{\text{stay}} - l_i^{\text{switch}} \right). \tag{3.5}$$

As long as there is no displacement then $l_i^{\text{stay}} = d$ and by revealed preference $l_i^{\text{switch}} \leq d$, so that $l_i \leq d$. Thus, switching enables workers to mitigate the losses from occupational decline. In the following subsections we verify that, in each version of our model, $\frac{\partial p_i}{\partial d} \geq 0, \frac{\partial l_i}{\partial d} \geq 0$ (with strict inequalities for some i): the larger the drop in demand, the more workers switch, and the higher are earnings losses. Furthermore, $\frac{\partial l_i}{\partial \alpha_{iA}} = -\frac{\partial p_i}{\partial \alpha_{iA}} \left(l_i^{\text{stay}} - l_i^{\text{switch}} \right) + p_i \frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}}$. In other words, losses decrease with initial within-occupation earnings rank if the switching probability is increasing and the loss of switchers decreasing in initial earnings rank, $\frac{\partial p_i}{\partial \alpha_{iA}} > 0$ and $\frac{\partial l_i^{\text{switch}}}{\partial \alpha_{iA}} < 0$.

In what follows, we investigate how mean earnings losses vary with α_{iA}, and hence with initial earnings, under various assumptions about switching costs and anticipation of the price change. To characterise switching behavior and earnings losses, we require a distributional assumption. For simplicity, we henceforth assume that α_{iA} and α_{iB} are independent and both uniformly distributed between zero and some finite but possibly large number $\overline{\alpha}$. We argue below that our results are robust to alternative distributional assumptions.

Subsection 3.A.2 Baseline model

We start with the simplest case, where occupational prices π_{kt} are revealed at the start of each period and there are no switching costs. Hence, occupational choice is a sequence of static decisions that can be analysed in isolation. The set of workers choosing occupation A in period 1 is characterised by the inequality $\alpha_{iB} \leq \alpha_{iA}$, and it lies on and below the main diagonal in panel (a) of Figure 3.3 (blue and red areas). The workers who switch in the second period must satisfy the inequalities $\alpha_{iB} \leq \alpha_{iA}$ and $\alpha_{iB} > \alpha_{iA} - d$, indicated by the blue area in panel (a) of Figure 3.3.

47. Formally,

$$\begin{aligned}
l_i &\equiv l_i(\alpha_{iA}, d) &&\equiv \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 1], \\
l_i^{\text{switch}} &\equiv l_i^{\text{switch}}(\alpha_{iA}, d) &&\equiv \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, k_{i2} = B, D_A = 1], \\
l_i^{\text{stay}} &\equiv l_i^{\text{stay}}(\alpha_{iA}, d) &&\equiv \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, D_A = 0] - \mathbb{E}[y_{i2}|\alpha_{iA}, k_{i1} = A, k_{i2} = A, D_A = 1], \\
p_i &\equiv p_i(\alpha_{iA}, d) &&\equiv \mathbb{P}(k_{i2} = B|k_{i1} = A, \alpha_{iA}, D_A = 1).
\end{aligned}$$

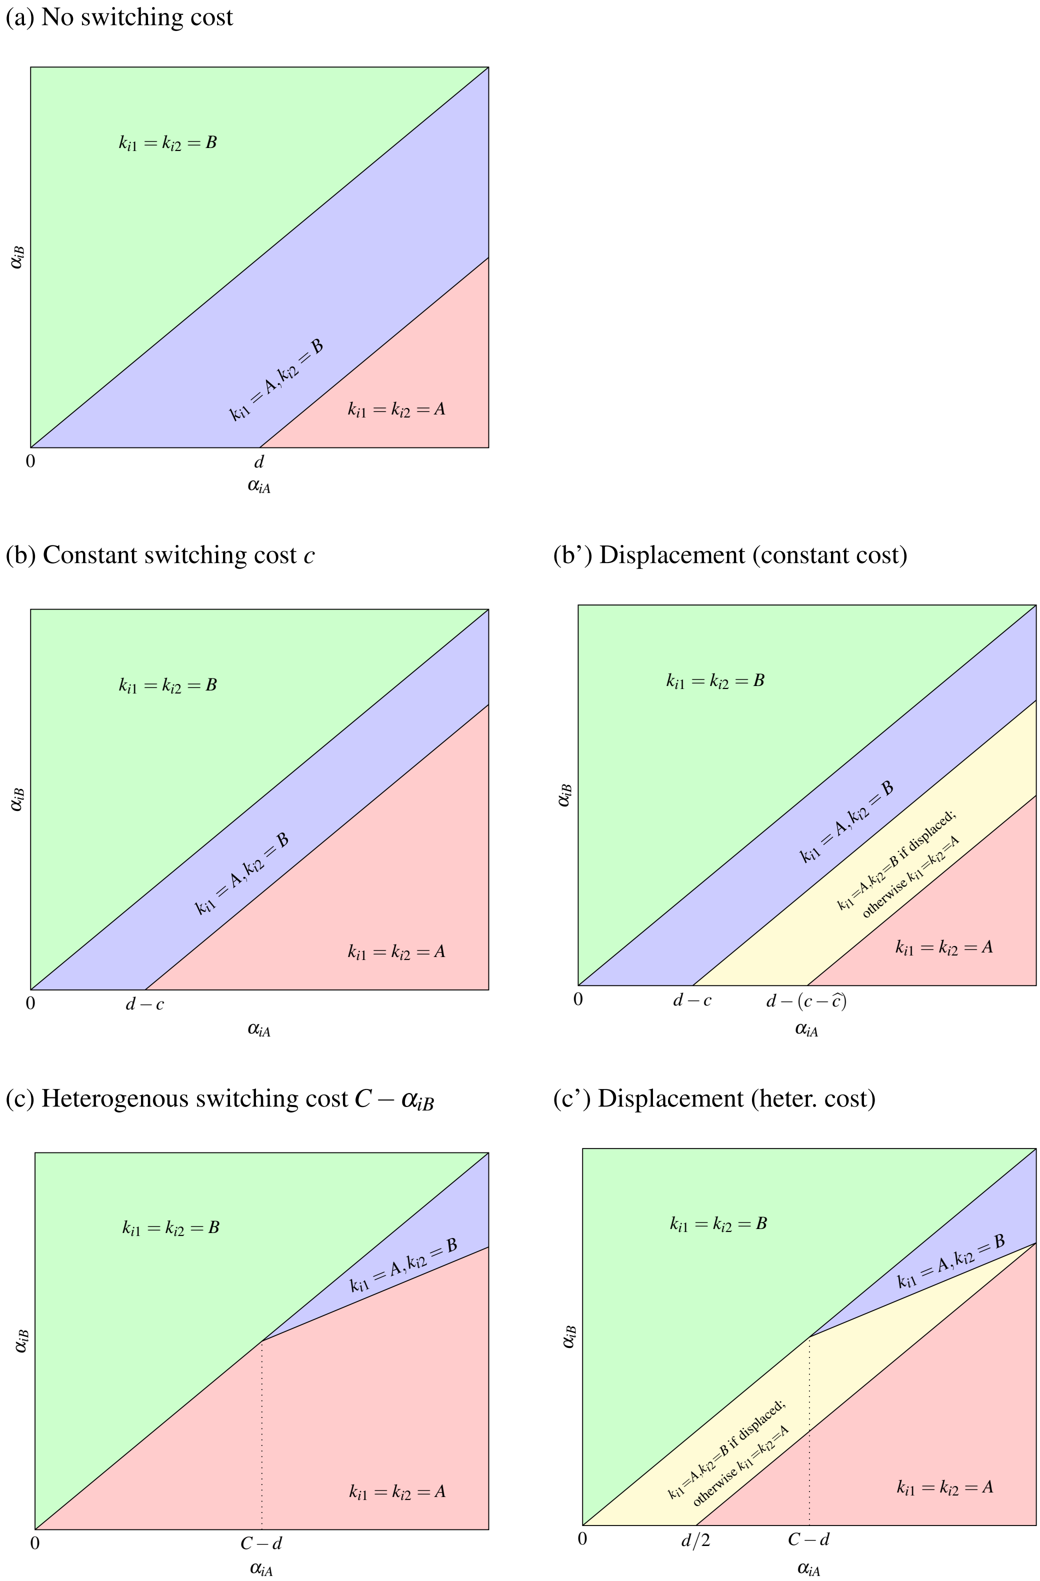

Notes: k_{it} denotes the occupation chosen by worker i in period t. α_{ik} denotes log productivity of worker i in occupation k. d is the amount by which the relative occupational log price declines from period 1 to period 2. The parameter values chosen are $(\overline{\alpha}, d, c, \widehat{c}, C) = (1, 0.5, 0.25, 0.25, 1)$.

Figure 3.3: Sorting in a two-period Roy model

To characterise switching probabilities and earnings losses, we need to distinguish two cases. Among workers in occupation A with $\alpha_{iA} \leq d$, everyone switches and their period-2 log earnings, given uniformity, are on average $\alpha_{iA}/2$, which is also the earnings loss they suffer. For those with $\alpha_{iA} > d$, the probability of switching is d/α_{iA}. The switchers' log productivity in occupation B lies between $\alpha_{iA} - d$ and α_{iA}, so given uniformity their period-2 log earnings are on average $\alpha_{iA} - d/2$, so that they suffer a loss of $d/2$. Switching probabilities, and their derivatives with respect to initial skill, are thus

$$p_i = \begin{cases} 1 & \text{if } \alpha_{iA} \leq d \\ \frac{d}{\alpha_{iA}} & \text{if } \alpha_{iA} > d, \end{cases} \qquad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq d \\ -\frac{d}{\alpha_{iA}^2} < 0 & \text{if } \alpha_{iA} > d, \end{cases}$$

and earnings losses are

$$l_i = \begin{cases} \frac{\alpha_{iA}}{2} & \text{if } \alpha_{iA} \leq d \\ d\left(1 - \dfrac{d}{2\alpha_{iA}}\right) & \text{if } \alpha_{iA} > d, \end{cases} \qquad \frac{\partial l_i}{\partial \alpha_{iA}} = \begin{cases} \frac{1}{2} > 0 & \text{if } \alpha_{iA} \leq d \\ \dfrac{d^2}{2\alpha_{iA}^2} > 0 & \text{if } \alpha_{iA} > d. \end{cases}$$

Given the above expressions, it is also straightforward to verify that $\partial p_i/\partial d \geq 0$ and $\partial l_i/\partial d \geq 0$. We summarise our analytical results for the baseline model as follows.

Result 1. *The probability of switching among those initially working in A is weakly <u>decreasing</u> in initial earnings. Moreover, switchers' earnings losses are also weakly increasing in initial earnings; and taken together, mean earnings losses for workers starting out in occupation A are strictly <u>increasing</u> in initial earnings.*

To understand the intuition for these results, call occupation A "typist" and occupation B "cashier", where typists suffer a negative demand shock. The worst typists could only become the worst cashiers, otherwise they would have chosen to be cashiers in period 1. But the best typists can at most become the best cashiers, and in general they will not all be the best cashiers. Therefore, the best typists are less able to mitigate their earnings losses by becoming cashiers, and they suffer larger losses than the worst typists.

This argument suggests that switching probabilities are decreasing and earnings losses are increasing in ability under a large set of alternative assumptions on the skill distributions. A sufficient condition for earnings losses to be higher for the most able than for the least able is that there is finite support with positive probability mass for all $(\alpha_{iA}, \alpha_{iB}) \in [0, \overline{\alpha}] \times [0, \overline{\alpha}]$.

Subsection 3.A.3 Costs of switching between occupations

We continue to assume that the period-2 price change is unanticipated, but now we assume that there are costs of switching occupations. We think of these costs as the time lost searching for a new job or spent in retraining, and model them as additive in log terms. We start with the simple case where the time cost is constant across workers (and thus proportional to earnings), and then consider a case where it is decreasing in workers' ability in the destination occupation. In both cases, we assume that the switching cost is symmetric between occupations.

Occupational choice is no longer a period-by-period decision. Instead, workers choose in period 1 the occupation with the highest expected present discounted value of log earnings, net of switching costs. Let us assume that occupational log prices follow a random walk, $\mathbb{E}[\widetilde{\pi}_2] = \widetilde{\pi}_1 = 0$, where the last equality is due to our earlier simplifying assumption.[48] This means that expected payoffs are equal across occupations apart from a worker's skills. Therefore, in both cases considered in this section we have that worker i chooses occupation A in period 1 if and only if $\alpha_{iA} \geq \alpha_{iB}$.

Constant switching costs

Take first the case where the cost for moving occupations is a constant $c \in (0, d)$; the case $c \geq d$ is uninteresting since nobody would switch in response to the adverse shock, so we only consider the case $c < d$.

Decisions at the beginning of the terminal period 2 are easily characterised, as before. After the price change, worker i switches if and only if $\alpha_{iB} - c > \alpha_{iA} - d$. Given period-1 choices, the workers who switch to occupation B after the price change satisfy the inequalities $\alpha_{iB} \leq \alpha_{iA}$ and $\alpha_{iB} > \alpha_{iA} - (d-c)$. Panel (b) of Figure 3.3 shows a situation that is qualitatively similar to the baseline model, except that the blue region marking the workers who switch is smaller than in panel (a). Deriving expressions for the switching probability and earnings loss as a function of initial earnings follows along very similar lines as in the proof of Result 1.

To characterise switching probabilities and earnings losses, we again to distinguish two cases. Among workers in occupation A with $\alpha_{iA} \leq d - c$, everyone switches and their period-2 log earnings, given uniformity, are on average $\alpha_{iA}/2 - c$, which gives an average loss of $\alpha_{iA}/2 + c$. For those with $\alpha_{iA} > d - c$, the probability of switching is $(d-c)/\alpha_{iA}$. The switchers' log productivity in occupation B lies between $\alpha_{iA} - (d-c)$ and α_{iA}, so given uniformity their period-2 log earnings are on average $\alpha_{iA} - (d-c)/2$, so that they suffer a loss of $(d+c)/2$. Switching probabilities, and their derivatives with respect to initial skill, are thus

$$p_i = \begin{cases} 1 & \text{if } \alpha_{iA} \leq d-c \\ \frac{d-c}{\alpha_{iA}} & \text{if } \alpha_{iA} > d-c, \end{cases} \qquad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq d-c \\ -\frac{d-c}{\alpha_{iA}^2} < 0 & \text{if } \alpha_{iA} > d-c, \end{cases}$$

and earnings losses are

$$l_i = \begin{cases} \frac{\alpha_{iA}}{2} + c & \text{if } \alpha_{iA} \leq d-c \\ d - \dfrac{(d-c)^2}{2\alpha_{iA}} & \text{if } \alpha_{iA} > d-c, \end{cases} \qquad \frac{\partial l_i}{\partial \alpha_{iA}} = \begin{cases} \frac{1}{2} > 0 & \text{if } \alpha_{iA} \leq d-c \\ \dfrac{(d-c)^2}{2\alpha_{iA}^2} > 0 & \text{if } \alpha_{iA} > d-c. \end{cases}$$

48. Instead of the random walk assumption we could impose that demand changes are somehow otherwise perfectly unforeseen, for instance due to adaptive expectations (in Section 3.A.5 we consider the case where demand changes are anticipated).

Given the above expressions, as in the baseline model, $\partial p_i/\partial d \geq 0$ and $\partial l_i/\partial d \geq 0$. We summarise our analytical results for the constant switching cost model as follows.

Result 2. *Under a constant switching cost, we obtain the same qualitative conclusions as in Result 1: The probability of switching among those initially working in A is weakly* <u>*decreasing*</u>*, switchers' earnings losses are weakly increasing, and mean earnings losses for workers starting out in occupation A are strictly* <u>*increasing*</u>*, in initial earnings.*

The same intuition as in the baseline model of Section 3.A.2 applies: the best workers in the declining occupation are less likely to be able to mitigate their earnings losses by switching occupation.

Heterogenous switching costs

Suppose instead that workers who wish to switch from A to B must pay a switching cost equal to $C-\alpha_{iB}$ and that workers who switch from B to A incur a cost $C-\alpha_{iA}$, with $C>\overline{\alpha}$ (the condition $C>\overline{\alpha}$ ensures that all workers face a strictly positive switching cost). This structure of switching costs captures in a reduced form way the frictions that occupational moves may entail: for example, job search may take time, and those more able in the new occupation may find a job more quickly.

After the price change in period 2, worker i switches if and only if $\alpha_{iB}-(C-\alpha_{iB})>\alpha_{iA}-d$. Thus, the workers who switch to occupation B after the shock must now satisfy the inequalities $\alpha_{iB}\leq\alpha_{iA}$ and $\alpha_{iB}>\alpha_{iA}/2+(C-d)/2$, shown as the blue area in panel (c) of Figure 3.3. The figure shows that workers with α_{iA} below $C-d$ do not switch, and that above $C-d$, the fraction switching is increasing in α_{iA} due to uniformity. Thus,

$$p_i=\begin{cases}0 & \text{if } \alpha_{iA}<C-d\\ \frac{1}{2}-\frac{C-d}{2\alpha_{iA}} & \text{if } \alpha_{iA}\geq C-d,\end{cases}\qquad \frac{\partial p_i}{\partial\alpha_{iA}}=\begin{cases}0 & \text{if } \alpha_{iA}<C-d\\ \frac{C-d}{2\alpha_{iA}^2}>0 & \text{if } \alpha_{iA}\geq C-d,\end{cases}$$

and

$$l_i^{\text{switch}}=\frac{d}{2}+\frac{C-\alpha_{iA}}{2},\qquad \frac{\partial l_i^{\text{switch}}}{\partial\alpha_{iA}}=-\frac{1}{2}<0$$

where we used the fact that mean earnings of switchers equal $3\alpha_{iA}/2-(C+d)/2$. Thus, if $\alpha_{iA}>C-d$, we have by (3.5) that $\partial l_i/\partial\alpha_{iA}<0$ (and zero otherwise). It is also straightforward to verify that $\partial p_i/\partial d\geq 0, \partial l_i/\partial d>0$. To summarise:

Result 3. *If the cost of switching occupations is decreasing in initial earnings, then the probability of switching among those initially working in A is weakly* <u>*increasing*</u> *in initial earnings, and mean losses conditional are (weakly)* <u>*decreasing*</u> *in initial earnings.*

In terms of the example above, in this case the worst typists do not switch, because their initial choice of occupation A reveals not only low earnings potential in occupation B but also a large switching cost. Among the best typists, however, many possess substantial earnings potential as cashiers, as well as low switching costs. Therefore, the best typists are on average better able to mitigate their earnings losses by becoming cashiers, and hence the earnings losses from the demand shock are smaller for the best typists than for the worst typists.[49]

Subsection 3.A.4 Job displacement

So far, we have been concerned with earnings losses as a function of initial earnings in the context of a Roy model where any moves between occupations are voluntary. By revealed preference, losses of movers must be less than those of stayers. Here we show that introducing job displacement and a cost of finding a new job in the initial occupation may overturn this result.[50]

Suppose that workers in either occupation experience job displacement with probability λ at the end of period 1. Occupations are still symmetric in expectation, so that period-1 choices are as before.

Displacement under constant switching costs

Displacement affects choices only in the presence of switching costs. First we assume that displaced workers incur a cost $\widehat{c} > 0$ to find a job in their starting occupation, and a cost c to find a job in the other occupation (the latter of course also applies to non-displaced workers). Here we have in mind exogenous job losses, for instance due to plant closure, which are a standard feature of search models (see for instance Pissarides 2000). To ensure that displaced workers in occupation B do not wish to switch to occupation A we assume that $\widehat{c} < c + d$.

The workers who are displaced switch occupation if and only if $\alpha_{iB} > \alpha_{iA} - (d - (c - \widehat{c}))$, and among them are individuals who would remain if not displaced, $\alpha_{iB} \leq \alpha_{iA} - (d - c)$. Workers not suffering displacement switch voluntarily if and only if $\alpha_{iB} > \alpha_{iA} - (d - c)$. Thus, there is a set of workers who switch occupation only if suffering displacement, as illustrated by the yellow area in panel (b') of Figure 3.1. Given uniformity, we have switching probabilities

$$p_i = \begin{cases} 1 & \text{if } \alpha_{iA} \leq d - c \\ \lambda + (1-\lambda)\frac{d-c}{\alpha_{iA}} & \text{if } d - c < \alpha_{iA} < d - (c - \widehat{c}) \\ \frac{d-c}{\alpha_{iA}} + \lambda \frac{\widehat{c}}{\alpha_{iA}} & \text{if } \alpha_{iA} \geq d - (c - \widehat{c}), \end{cases} \qquad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq d - c \\ -(1-\lambda)\frac{d-c}{\alpha_{iA}^2} < 0 & \text{if } d - c < \alpha_{iA} < d - (c - \widehat{c}) \\ -\frac{d-c}{\alpha_{iA}^2} - \lambda \frac{\widehat{c}}{\alpha_{iA}^2} < 0 & \text{if } \alpha_{iA} \geq d - (c - \widehat{c}). \end{cases}$$

The earnings loss in the first region is the same as in the case without displacement in Section 3.A.3. The calculations are more involved in the second and third regions. Let ρ_i^{blue} denote the probability that a worker's skill α_{iB} lies in the blue region in panel (b') of Figure 3.3 (conditional on starting out in occupation A and on α_{iA}), let l_i^{blue} be her expected loss, and analogously define $\rho_i^{\text{yellow}}, \rho_i^{\text{red}}$ and

49. While our model excludes occupation-specific human capital, it does allow us to think about some of its potential implications. For example, if all workers accumulate occupation-specific human capital additively (in logarithms) the effects are similar to adding constant switching costs, since switching means foregoing this capital. And if the accumulation of occupation-specific human capital is faster for those with higher ability in the occupations they initially select, they become less mobile, in contrast to the case of heterogeneous switching costs discussed above. Either way, adding occupation-specific human capital does not help to rationalise our empirical findings.

50. Recall that a large literature has documented substantial earnings losses due to job displacement (see for instance Jacobson, LaLonde, and Sullivan 1993) and even larger losses if such displacement coincides with switching occupation (Kambourov and Manovskii 2009).

$l_i^{\text{yellow}}, l_i^{\text{red}}$. When $d-c < \alpha_{iA} < d-(c-\widehat{c})$, we have

$$\left(\rho_i^{\text{yellow}}, \rho_i^{\text{blue}}\right) = \left(1 - \frac{d-c}{\alpha_{iA}}, \frac{d-c}{\alpha_{iA}}\right), \qquad \left(l_i^{\text{yellow}}, l_i^{\text{blue}}\right) = \left(\lambda\left(\frac{\alpha_{iA}}{2} + \frac{d+c}{2}\right) + (1-\lambda)d, \frac{d+c}{2}\right)$$

and

$$l_i = l_i^{\text{blue}} + \rho_i^{\text{yellow}}\left(l_i^{\text{yellow}} - l_i^{\text{blue}}\right),$$

where we note that $l_i^{\text{yellow}} - l_i^{\text{blue}} = \lambda\frac{\alpha_{iA}}{2} + (1-\lambda)\frac{d-c}{2} > 0$.

When $\alpha_{iA} \geq d-(c-\widehat{c})$, we additionally define ρ_i^{red} and l_i^{red} in the same sense as above, so we have

$$\left(\rho_i^{\text{red}}, \rho_i^{\text{yellow}}, \rho_i^{\text{blue}}\right) = \left(1 - \frac{d-(c-\widehat{c})}{\alpha_{iA}}, \frac{\widehat{c}}{\alpha_{iA}}, \frac{d-c}{\alpha_{iA}}\right), \qquad \left(l_i^{\text{red}}, l_i^{\text{yellow}}, l_i^{\text{blue}}\right) = \left(\lambda(d+\widehat{c}) + (1-\lambda)d, d + \frac{\widehat{c}}{2}, \frac{d+c}{2}\right)$$

and

$$l_i = l_i^{\text{blue}} + \rho_i^{\text{red}}\left(l_i^{\text{red}} - l_i^{\text{blue}}\right) + \rho_i^{\text{yellow}}\left(l_i^{\text{yellow}} - l_i^{\text{blue}}\right),$$

where $l_i^{\text{red}} - l_i^{\text{blue}} = \frac{d-c}{2} + \lambda\widehat{c} > 0$ and $l_i^{\text{yellow}} - l_i^{\text{blue}} = \frac{d-c}{2} + \frac{\widehat{c}}{2} > 0$. Thus,

$$l_i = \begin{cases} \frac{\alpha_{iA}}{2} + c & \text{if } \alpha_{iA} \leq d-c \\ \frac{d+c}{2} + \left(1 - \frac{d-c}{\alpha_{iA}}\right)\left(\lambda\frac{\alpha_{iA}}{2} + (1-\lambda)\frac{d-c}{2}\right) & \text{if } d-c < \alpha_{iA} < d-(c-\widehat{c}) \\ d + \lambda\widehat{c} - \frac{(d-(c-\widehat{c}))(d-c+2\lambda\widehat{c})}{2\alpha_{iA}} & \text{if } \alpha_{iA} \geq d-(c-\widehat{c}), \end{cases}$$

and

$$\frac{\partial l_i}{\partial \alpha_{iA}} = \begin{cases} \frac{1}{2} > 0 & \text{if } \alpha_{iA} \leq d-c \\ \frac{d-c}{\alpha_{iA}^2}\left(l_i^{\text{yellow}} - l_i^{\text{blue}}\right) + \rho_i^{\text{yellow}}\frac{\lambda}{2} > 0 & \text{if } d-c < \alpha_{iA} < d-(c-\widehat{c}) \\ \frac{(d-(c-\widehat{c}))(d-c+2\lambda\widehat{c})}{2\alpha_{iA}^2} > 0 & \text{if } \alpha_{iA} \geq d-(c-\widehat{c}). \end{cases}$$

It is also straightforward to verify that $\partial p_i/\partial d \geq 0$ and $\partial l_i/\partial d \geq 0$. To summarise:

Result 4. *If the cost of switching occupations is constant, if workers may be displaced from their jobs, and if the cost of finding a new job in the starting occupation is also constant, then the probability of switching among those initially working in A is (weakly) decreasing in initial earnings, and mean losses are strictly increasing in initial earnings.*

Displacement does not affect the qualitative results that we obtained before when assuming constant switching costs (Section 3.A.3). It is still the case that the most skilled among the workers starting out in A are less likely to be in a position where they can mitigate their earnings losses by switching occupation.

Displacement under heterogenous switching costs

In the case of a heterogeneous cost of switching occupations, we introduce in symmetric fashion a cost for finding a job in the starting occupation, if displaced. A worker starting out in occupation A incurs a cost $C-\alpha_{iB}$ for moving to occupation B, and if displaced incurs a cost $C-\alpha_{iA}$ for finding another job in occupation A. For a worker starting in B, the costs are $C-\alpha_{iA}$ and $C-\alpha_{iB}$, respectively.

Period-1 occupational choices are as before due to symmetry. Recall that workers not affected by displacement switch voluntarily at the start of period 2 if and only if $\alpha_{iB} > \alpha_{iA}/2 + (C-d)/2$. Workers that do suffer displacement switch occupation if and only if $\alpha_{iB} > \alpha_{iA} - d/2$, as illustrated by the yellow area in (c') of Figure 3.1. The figures suggests that, going from low to high values of α_{iA}, the switching probability first equals λ as all displaced workers switch and then falls below λ and becomes a decreasing function of α_{iA}, and possibly, eventually an increasing function of α_{iA}, since the incidence of voluntary switching increases in α_{iA} for large values of α_{iA}.

Indeed, switching probabilities are characterised as

$$p_i = \begin{cases} \lambda & \text{if } \alpha_{iA} \leq \frac{d}{2} \\ \frac{\lambda d}{2\alpha_{iA}} & \text{if } \frac{d}{2} < \alpha_{iA} < C-d \\ \frac{1-\lambda}{2} - \frac{(1-\lambda)C-d}{2\alpha_{iA}} & \text{if } \alpha_{iA} \geq C-d, \end{cases} \qquad \frac{\partial p_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq \frac{d}{2} \\ -\frac{\lambda d}{2\alpha_{iA}^2} < 0 & \text{if } \frac{d}{2} < \alpha_{iA} < C-d \\ \frac{(1-\lambda)C-d}{2\alpha_{iA}^2} \gtreqless 0 \Leftrightarrow (1-\lambda)C \gtreqless d & \text{if } \alpha_{iA} \geq C-d. \end{cases}$$

The expression for the switching probability for $\alpha_{iA} \geq C-d$ follows from the fact that when $\alpha_{iA} \geq C-d$, the probability of a worker being in the blue region in panel (c') of Figure 3.3 equals $1/2 - (C-d)/(2\alpha_{iA})$, and the probability of being in the yellow region is $C/(2\alpha_{iA}) - 1/2$. We see that for $\alpha_{iA} > C-d$ the switching probability is strictly increasing in α_{iA} provided the shock is not too large, $d/C < 1-\lambda$.

The losses from occupational decline in this version of the model are

$$l_i = \begin{cases} \lambda C + (1-\lambda)d & \text{if } \alpha_{iA} \leq \frac{d}{2} \\ \lambda \left[C - \alpha_{iA} + d - \frac{d^2}{4\alpha_{iA}}\right] + (1-\lambda)d & \text{if } \frac{d}{2} < \alpha_{iA} < C-d \\ \ell_i & \text{if } \alpha_{iA} \geq C-d, \end{cases}$$

where ℓ_i is to be characterised below. For $\alpha_{iA} < d/2$, workers move if and only if they are displaced, and the expected loss of movers is constant at C, while the expected loss of stayers is constant at d. Hence, the overall expected loss in this region equals $\lambda C + (1-\lambda)d$ and thus does not depend on α_{iA}. For $\alpha_{iA} \in (d/2, C-d)$, the expected loss for those who are displaced and stay is $C - \alpha_{iA} + d$ and for those who are displaced and move it is $C - \alpha_{iA} + d/2$ (all non-displaced workers stay). Together with the switching probability, the result follows.

Finally, we consider the case $\alpha_{iA} \geq C-d$. Let ρ_i^{red} denote the probability that a worker's skill α_{iB} lies in the red region in panel (c') of Figure 3.3 (conditional on starting out in occupation A and on α_{iA}), let l_i^{red} be her expected loss, and analogously define $\rho_i^{\text{yellow}}, \rho_i^{\text{blue}}$ and $l_i^{\text{yellow}}, l_i^{\text{blue}}$. We have

$$\ell_i = l_i^{\text{blue}} + \rho_i^{\text{red}} \left(l_i^{\text{red}} - l_i^{\text{blue}} \right) + \rho_i^{\text{yellow}} \left(l_i^{\text{yellow}} - l_i^{\text{blue}} \right),$$

where

$$\left(\rho_i^{\text{red}}, \rho_i^{\text{yellow}}\right) = \left(1 - \frac{d}{2\alpha_{iA}}, \frac{C-\alpha_{iA}}{2\alpha_{iA}}\right), \qquad \left(l_i^{\text{red}}, l_i^{\text{yellow}}, l_i^{\text{blue}}\right) = \left(d + \lambda(C-\alpha_{iA}), d + \lambda\frac{C-\alpha_{iA}}{2}, \frac{d}{2} + \frac{C-\alpha_{iA}}{2}\right),$$

and so

$$\ell_i = d + \lambda(C-\alpha_{iA}) + \frac{1}{2\alpha_{iA}}\left(-\frac{d^2}{2} + (1-\lambda)(C-\alpha_{iA})\left(d - \frac{C-\alpha_{iA}}{2}\right)\right)$$

and

$$\frac{\partial \ell_i}{\partial \alpha_{iA}} = \frac{d^2 + C(C-2d)(1-\lambda) - \alpha_{iA}^2(3\lambda+1)}{4\alpha_{iA}}.$$

Thus,[51]

$$\frac{\partial l_i}{\partial \alpha_{iA}} = \begin{cases} 0 & \text{if } \alpha_{iA} \leq \frac{d}{2} \\ \lambda\left[\frac{d^2}{4\alpha_{iA}^2} - 1\right] < 0 & \text{if } \frac{d}{2} < \alpha_{iA} < C-d \\ \partial \ell_i/\partial\alpha_{iA} < 0 \Leftarrow \frac{d}{C} < \frac{2}{3} & \text{if } \alpha_{iA} \geq C-d. \end{cases}$$

Let us assume that $d/C < \min\{2/3, 1-\lambda\}$. This is a sensible assumptions, as it implies that the shock is not huge—it does not come close to making the declining occupation vanish (recall that $C > \overline{\alpha}$). This also is the case that gives rise to the richest patterns of switching behavior, as for instance the yellow region of panel (c') of Figure 3.3 would not coexist with the blue region if $d/2 > C-d$. As before, $\partial p_i/\partial d \geq 0$ and $\partial l_i/\partial d \geq 0$. To summarise:

Result 5. *If the cost of switching occupations is decreasing in initial earnings, if workers may be displaced from their jobs, and if the cost of finding a new job in the starting occupation is also decreasing in initial earnings, then the probability of switching among those initially working in A is U-shaped in initial earnings, and mean losses are (weakly) decreasing in initial earnings.*

Intuitively, the earnings loss is decreasing in initial earnings, as in the case with heterogenous switching costs without displacement (Section 3.A.3), because the costs of moving jobs—both within and across occupations—decrease with initial earnings.

Subsection 3.A.5 Revelation of period-2 prices at the start of period 1

As a final variation on our model, we consider a case where period-2 prices are revealed to be $\widetilde{\pi}_2 = d$ at the start of period 1. Without switching costs, decisions are again static and occupational choices follow

51. The sign of $\frac{\partial \ell_i}{\partial \alpha_{iA}}$ is the same as the sign of its numerator. Since $\alpha_{iA} > C-d$ when $l_i = \ell_i$, the numerator is strictly less than

$$d^2 + C(C-2d)(1-\lambda) - (C-d)^2(3\lambda+1) = \lambda\left(d^2 - 4(C-d)^2\right)$$

which is negative if $d/C < 2/3$. The expression is also negative if $d/C > 2$, but cases with $d > C$ are uninteresting since they imply that everyone leaves the declining occupation.

the same conditions as in the baseline model of Section 3.A.2. Suppose however that there is a constant (across individuals) switching cost $c \in (0,d)$ for moving between occupations, as in the first scenario considered in Section 3.A.3. Workers choose an occupational path by comparing the deterministic life-time utilities associated with the choices (A,A), (A,B), and (B,B). The life-time utilities are given by $V_{iAA} = \alpha_{iA} + \beta(\alpha_{iA} - d)$, $V_{iAB} = \alpha_{iA} + \beta(\alpha_{iB} - c)$, and $V_{iBB} = \alpha_{iB} + \beta\alpha_{iB}$. First, let us assume that switching cots are not too large, $(1+\beta)c < d$. Then we have:

- If $\alpha_{iB} \leq \alpha_{iA} - (d-c)$, the worker chooses (A,A).
- If $\alpha_{iB} > \alpha_{iA} - (d-c)$ and $\alpha_{iB} \leq \alpha_{iA} - \beta c$, the worker chooses (A,B).
- If $\alpha_{iB} > \alpha_{iA} - \beta c$, the worker chooses (B,B).

All workers with $\alpha_{iB} > \alpha_{iA} - \beta c$ choose occupation B in period 1 and remain there. Thus, some workers who otherwise would have started out in occupation A instead start in B to avoid the switching cost, and the fraction of workers switching in the period when the shock hits is smaller than without anticipation of the shock.

If switching costs are large instead, $(1+\beta)c \geq d$, then workers with $\alpha_{iB} \leq \alpha_{iA} - \beta c$ choose (A,A) and workers with $\alpha_{iB} > \alpha_{iA} - \beta c$ choose (B,B), so that no switching occurs after period 1. To summarise:

Result 6. *If period-2 prices are revealed already at the start of period 1, and under a constant occupational switching cost, the fraction of workers starting out in occupation A, and the fraction of workers leaving occupation A after the first period, are both smaller than in the case without anticipation discussed in Section 3.A.3. The fraction switching occupation after period 1 may even be zero if the switching cost is large.*

More generally, the model suggests that the set of workers who are in declining occupations may differ for anticipated and unanticipated shocks. Different combinations of anticipation, general equilibrium responses, heterogeneity of occupational switching costs, and displacement, may lead to a range of different outcomes.

Subsection 3.A.6 Summary of theoretical results

We have modelled occupational decline using a Roy model, where employment in an occupation declines as a result of a fall in occupational price caused by a technology shock. The model illustrates how earnings losses due to occupational decline are mitigated by occupational switching.

Furthermore, our frictionless baseline model makes three predictions: the probability of leaving a declining occupation is decreasing in initial earnings; earnings losses due to occupational decline are increasing in initial earnings; and earnings losses of those who leave a declining occupation are less than the losses of those who remain.

Anticipating that these predictions are inconsistent with our empirical findings, we have considered several modifications to the model. Introducing an occupational switching cost that is decreasing in the worker's earnings in the destination occupation leads to a positive relationship between switching probabilities and initial earnings, and a negative relationship between earnings losses and initial earnings. Allowing for displacement, together with a cost of switching jobs within an occupation, implies that switchers' earnings losses may be larger than those of stayers. Moreover, displacement can cause switching probabilities to be U-shaped in initial earnings, whereby low-earning workers switch involuntarily if displaced, while high-earning workers switch voluntarily regardless of displacement.

The importance of switching costs in our theoretical analysis motivates our empirical approach of focusing not only on losses in career earnings incurred by workers starting out in declining occupations, but also on losses in years employed, as well as on the incidence of unemployment and retraining. While our model does not include a non-work sector, it could be shown that a negative demand shock would trigger moves from the affected occupation into non-participation.

Finally, we have used our model to show that much of re-sorting in response to a technology shock may occur before the shock hits if it is anticipated in advance, motivating our investigation of both anticipated and unanticipated occupational decline.

Appendix 3.B Data appendix

Subsection 3.B.1 Data sources

Our main analysis is based on individual-level longitudinal administrative data covering the entire population of Sweden 1985-2013, and on various editions of the Occupational Outlook Handbook (OOH) published by the Bureau of Labour Statistics (BLS) in the US.

Occupation data

Our primary source for measuring occupational decline are the 1986-87 and the 2018-19 editions of the Occupational Outlook Handbook (Bureau of Labor Statistics 1986, 2018). The OOH contains a description of the nature of work, the current number of jobs, and projected employment growth for hundreds of occupations. For a subset of these occupations, more detailed information is reported, including required qualifications, pay, and the role of technology: whether technology is expected to affect—or has already affected—the occupation in question, and if so, what the impact on employment will be or has been. In the 1986-87 edition, 401 occupations are described, covering about 80% of US employment. Detailed information is available for 196 occupations, covering about 60% of employment.[52]

Swedish microdata

The main outcomes we study come from Swedish microdata. We obtain basic demographic (year of birth, gender, education, and place of residence) and labour market (employment status, annual earnings, industry, firm identifier) variables from the Integrated Database for Labour Market Research (LISA), a collection of administrative registers that is—like all our other Swedish data sources—provided by Statistics Sweden. During the period 1985-2013, LISA contains one observation per year on every individual aged 16-64 living in Sweden. Employment status and industry (as well as county of residence) are measured in November each year.

We also use individual-level data from the Swedish Public Employment Service (PES), which contain information on the total number of days registered with the PES, number of days registered as unemployed, and number of days spent in retraining programs administered by the PES, for all individuals ever registered with the PES during the years 1992-2013.

Our data on workers' occupations come from the population censuses, which were conducted every five years from 1960-1990, and from the Wage Structure Statistics (WSS) for the years 1996-2013.[53] The WSS contains the population of public sector workers and a sample of about 50% of private sector workers. Sampling is at the level of firms, and large firms are over-sampled. We apply sampling weights when working with the occupation variable from the WSS.

52. The number of distinct occupations in the OOH, as well as the number of occupations covered in detail, tends to increase over time. This means that our crosswalk from the 1986-87 to the 2018-19 edition is mostly, though not always, one-to-many.

53. We also obtain individual-level earnings data for 1975 and 1980 from the population censuses, which we use for falsification checks.

A useful feature of our data is that, in the 1985 and 1990 censuses, workers' occupation is coded using a 5-digit classification, YRKE5, containing about 1,400 distinct occupations. This allows us to accurately merge occupation-level information from the US (see below). Unfortunately, such detailed occupation codes are not available after 1990. From 1996-2013, a 3-digit classification containing 172 distinct codes, SSYK96, is available in the WSS. This classification is of a different nature than YRKE5, and the cross-walk between YRKE5 and SSYK96 likely introduces measurement error in workers' occupations after 1990.[54] This is an important caveat to our analysis of occupational employment shifts and individual workers' occupational mobility during 1985-2013.

Finally, adding the 1960 census allows us to calculate prior occupational employment changes at the 3-digit level using the YRKE3 classification, a coarser version of YRKE5 (there are 229 distinct codes that cover the period 1960-85).[55]

Subsection 3.B.2 Construction of variables

Occupation variables (OOH)

Using the reported employment numbers from our two editions of the OOH, we calculate the percentage growth in employment 1984-2016.[56] We manually map occupations across the two editions. If, after a careful search, a 1986-87 occupation has no counterpart in the 2018-19 edition, we classify it as having vanished, and assigned a percentage growth of -100.[57] While few occupations actually disappeared, examples of occupations that declined sharply include both white-collar occupations (typists, drafters, and telephone operators) and blue-collar ones (precision assemblers, welders, and butchers).

We also record for each US occupation its projected employment growth from the 1986-87 OOH. The BLS bases these predictions on (forecasts of) the size and demographic composition of the labour force, aggregate economic growth, commodity final demand, industry-level output and employment, the input-output matrix, and occupational employment and vacancies. The forecasts are not reported in percentage terms but grouped into the categories "declining", "little or no change", "increasing slower than average", "increasing about as fast as average", and "increasing faster than average". We create a cardinal predicted growth index assigning these categories the numbers 1-5 (where higher numbers correspond to more positive predicted employment changes). We report results both from using this index and using the categorical outlook variable.[58]

Merging of OOH variables to Swedish microdata, and defining occupational decline

In order to merge the OOH-based variables to Swedish data, we map the 401 1986-87 OOH occupations to the 1,396 5-digit Swedish occupation codes available in the 1985 census. We successfully map 379

54. Within broad types of jobs, SSYK96 also distinguishes occupations by the skill level of the workers.
55. The Swedish word *yrke* means occupation. SSYK stands for (the Swedish translation of) Swedish Standard Classification of Occupations.
56. The 1986-87 OOH reports employment numbers for 1984, while the 2018-19 edition reports 2016 employment figures.
57. Between the 1986-87 and 2018-19 editions of the OOH, some occupations were split or merged, which we take into account when computing the percentage growth. The details of this calculation are discussed later in this section.
58. Veneri (1997) uses US employment data to evaluate the ex-post accuracy of the projections used in the 1986-87 OOH, and concludes that they correctly foresaw most occupational trends, although there were also non-trivial sources of error.

US occupations to 1,094 Swedish occupations—we are able to find corresponding US occupations for 91% of Swedish workers in 1985. We map percentage changes in US employment 1984-2016, as well as 1986-87 OOH predictions (categorical and index), to Swedish 5-digit occupations using our crosswalk, applying weights (OOH 1984 employment shares) in the case of many-to-one matches.

We define a Swedish 5-digit occupation as declining if the weighted employment growth of its corresponding OOH occupations is negative and larger (in absolute magnitude) than 25%. We regard this as a sensible threshold: smaller declines may be the result of measurement error, as we had to exercise judgment in matching OOH occupations over time. At the same time, we report robustness checks using a number of alternative thresholds. We also use information from the OOH to determine whether technology likely played a role in the decline.[59] In 1985, 13% of Swedish employees worked in subsequently declining occupations, and 8% worked in subsequently declining occupations where the decline is likely linked to technology. We now provide more details on the process.

Assigning US OOH employment growth to Swedish occupations given a hypothetically unchanging OOH classification For clarity, we first describe what the calculation of employment growth would be if the OOH classification had not changed between the 1986-87 and 2018-19 editions. We then describe the adjustments we make given that the OOH classification did change.

The percentage change that we assign to each Swedish occupation s in the hypothetical case of an unchanging OOH classification is given by

$$g_s \equiv \frac{N_{s,2016} - N_{s,1984}}{N_{s,1984}}, \tag{3.6}$$

where $N_{s,t} \equiv \sum_{k \in \mathbb{K}_s} N_{k,t}$ is the sum of all year-t employment in the $k \in \mathbb{K}_s$ OOH occupations to which the Swedish YRKE5 occupation is matched. This percentage change can alternatively be expressed as

$$g_s \equiv \underset{1 \times K}{\alpha_s} \times \underset{K \times 1}{g}, \tag{3.7}$$

where the vector α_s is a vector of weights of length K, where K is the total number of OOH occupations in the 1986-87 OOH. Each element $\alpha_{s,k}$ represents the share of OOH occupation k in the mapping to Swedish occupation s, and it is based on the employment figures in the initial period 1984.[60] Thus,

59. To determine whether technology played a role in the decline, we proceeded as follows. We first applied the 25-percent cutoff to the OOH data to identify the declining occupations in the US. For the declining occupations we searched their detailed descriptions in the 1986-87 OOH for discussions of potential replacement of human labour by specific technologies, such as computers or robots. For the occupations lacking detailed descriptions in the 1986-87 OOH, we further searched one and two decades ahead, using the 1996-97 and 2006-07 editions (Bureau of Labor Statistics 1996, 2006), since in some cases occupations were re-grouped and so received detailed descriptions in those editions. Note that, while the OOH contains little backward-looking information on technology's role, it provides rich information on imminent technological changes expected to affect occupations. Conditional on an OOH occupation being classified as declining, we regard this information as reliable with respect to technology's role in the decline.

For those OOH occupations that we identified to have undergone technology-related declines, we map employment growth to Swedish 5-digit occupations creating a separate variable, technology-related employment growth. We define a Swedish 5-digit occupation as declining and linked to technology if the technology-related employment growth in the corresponding OOH occupations is below negative 25%. All technology-related declining occupations are declining occupations by construction, but some declining occupations may not be classified as having a technology link.

60. Note that the 1986-87 OOH uses data from 1984. Thus, the initial period is 1984 as far as US employment figures are concerned, but the data are extracted from a 1986 publication.

$\alpha_{s,k} \in [0,1]$, the vector α_s differs between Swedish YRKE5 occupations and its elements always sum to one. The vector g is filled with the 1984-2016 growth rates of all K OOH occupations. Formally,

$$\alpha_{s,k} \equiv \frac{\mathbb{1}_{k\in\mathbb{K}_s} \times N_{k,1984}}{\sum_{k\in\mathbb{K}_s} N_{k,1984}}, \qquad g_k \equiv \frac{N_{k,2016} - N_{k,1984}}{N_{k,1984}}.$$

The equivalence of (3.6) and (3.7) is easily shown:

$$\begin{aligned}
g_s &\equiv \frac{N_{s,2016} - N_{s,1984}}{N_{s,1984}} \\
&\equiv \frac{\sum_{k\in\mathbb{K}_s} N_{k,2016} - \sum_{k\in\mathbb{K}_s} N_{k,1984}}{\sum_{k\in\mathbb{K}_s} N_{k,1984}} \\
&\equiv \frac{\sum_{k\in\mathbb{K}_s} N_{k,1984} \times g_k}{\sum_{k\in\mathbb{K}_s} N_{k,1984}} \\
&\equiv \sum_k \frac{\mathbb{1}_{k\in\mathbb{K}_s} \times N_{k,1984} \times g_k}{\sum_{k\in\mathbb{K}_s} N_{k,1984}} \\
&\equiv \sum_k \alpha_{s,k} \times g_k \\
&\equiv \alpha_s \times g.
\end{aligned}$$

Assigning US OOH employment growth to Swedish occupations given the changing OOH classification The computation of the total changes in equation (3.6), or the weights and changes in equation (3.7) would be straightforward if the OOH occupation classification remained constant between the 1986-87 and 2018-19 editions. Alas, it did not, and so we need to adjust the calculation for any splits and merges that took place.

To see this, consider the following example: the OOH occupation "343 Metal pourers and casters, basic shapes" had employment 12,000 in 1984. By 2016, it had been merged with sixteen other occupations to "Metal and Plastic Machine Workers", with employment 1,039,600. It is obviously wrong to calculate the change in occupation "343 Metal pourers and casters, basic shapes" as a more than 85-fold increase:

$$g_{343} = \frac{1,039,600 - 12,000}{12,000} = 85.63$$

Instead, it is reasonable to sum the employment of all the seventeen merged occupations in 1984, with a total employment of 1,457,000, and calculate the change as

$$\hat{g}_{343} = \frac{1,039,600 - 1,457,000}{1,457,000} = -0.286$$

obtaining a 28.6% decline.

However, what happens to the weights in α_s? If we were to weight the "343 Metal pourers and casters, basic shapes" by their adjusted employment figure for 1984 (1,457,000), this occupation would seem 121 larger than it actually was (12,000). This creates problems when "343 Metal pourers and casters, basic shapes" is matched to a Swedish YRKE5 occupation that is also matched to other OOH occupations.

Consider, for instance, the Swedish YRKE5 occupation "732.50 Precision founder" to which "343 Metal pourers and casters, basic shapes" is matched, together with another OOH occupation "344 Molders and casters, hand".

Swedish YRKE5 occupation	OOH occupation	Employment in 1984	$\hat{g}_k$
732.50 Precision founder	343 Metal pourers and casters, basic shapes	12,000	-0.286
	344 Molders and casters, hand	17,000	-1.000

"344 Molders and casters, hand" was larger than "343 Metal pourers and casters, basic shapes" in 1984, and disappeared completely between 1984 and 2016. It seems like we should assign the Swedish YRKE5 occupation "732.50 Precision founders" a decline somewhere in between -28.6% and -100%, but closer to -100% since the disappearing occupation dominates. However, if we were to use *adjusted* employment figures when calculating the weights, "343 Metal pourers and casters, basic shapes" would be weighted as follows:

$$\hat{\alpha}_{s,343} = \frac{1,457,000}{1,457,000 + 17,000} = 0.988$$

That is, "343 Metal pourers and casters, basic shapes" would seem to account for almost *all* employment in the Swedish YRKE5 occupation, instead of less than half. This means that the weighted change will be mistakenly computed as

$$\hat{\alpha}_{s,343} \times \hat{g}_{343} + \alpha_{s,344} \times \hat{g}_{344}$$

$$= 0.988 \times (-0.286) + 0.012 \times (-1.00) = -0.295$$

Instead, we ought to use the original employment figures when calculating the weights. Then,

$$\alpha_{s,343} = \frac{12,000}{12,000 + 17,000} = 0.414$$

i.e. the OOH occupation "343 Metal pourers and casters, basic shapes" makes up 41.4% of employment in the Swedish YRKE5 occupation. Thus,

$$\alpha_{s,343} \times \hat{g}_{343} + \alpha_{s,344} \times \hat{g}_{344}$$

$$= 0.414 \times (-0.286) + 0.586 \times (-1.00) = -0.704$$

That is, the employment growth assigned to "732.50 Precision founders" should be -70.5%. We will thus treat weights and growth rates separately: The weights α_s are computed using the original employment figures, and the growth rates g_k are computed using the adjusted employment figures,

$$\hat{g}_s = \alpha_s \times \hat{g}. \tag{3.8}$$

The formal definition of our declining indicator is thus

$$\text{Declining}_s \equiv \mathbb{1}\{\alpha_s \times \hat{g} < -0.25\}.$$

It remains to specify how exactly the growth rates should be adjusted for splits and merges.[61]

- One-to-one: OOH occupations that were neither split or merged between the 1986-87 and 2018-19 editions of the OOH. No adjustment is needed, and the growth rate is defined as above,

$$\hat{g}_k = g_k \equiv \frac{N_{k,2016} - N_{k,1984}}{N_{k,1984}}.$$

- Many-to-one merge: Many 1984 occupations $k \in \mathbb{K}$ (where $\mathbb{K}$ is a set of 1984 occupations) were merged into one 2016 occupation $\tilde{k}$. 1984 employment figures of all merged occupations are summed and compared to the 2016 figures.

$$\hat{g}_{k\in\mathbb{K}} = \frac{N_{\tilde{k},2016} - \sum_{k\in\mathbb{K}} N_{k,1984}}{\sum_{k\in\mathbb{K}} N_{k,1984}}$$

- One-to-many split: One 1984 occupation k was split into many 2016 occupations $\tilde{k} \in \tilde{\mathbb{K}}$ (where $\tilde{\mathbb{K}}$ is a set of 2016 occupations). The 2016 employment figures of all resulting splits are added and compared to the 1984 figures.

$$\hat{g}_k = \frac{\sum_{\tilde{k}\in\tilde{\mathbb{K}}} N_{\tilde{k},2016} - N_{k,1984}}{N_{k,1984}}$$

- Many-to-many: Many 1984 occupations $k \in \mathbb{K}$ (where $\mathbb{K}$ is a set of 1984 occupations) were distributed into many 2016 occupations $\tilde{k} \in \tilde{\mathbb{K}}$ (where $\tilde{\mathbb{K}}$ is a set of 2016 occupations). The 1984 and 2016 employment figures are added and compared.

$$\hat{g}_{k\in\mathbb{K}} = \frac{\sum_{\tilde{k}\in\tilde{\mathbb{K}}} N_{\tilde{k},2016} - \sum_{k\in\mathbb{K}} N_{k,1984}}{\sum_{k\in\mathbb{K}} N_{k,1984}}$$

Identifying technology-related declines Having calculated the adjusted employment growth $\hat{g}_k$ for all occupations present in the 1986-87 OOH, we concentrate on those that declined sharply, $\hat{g}_k < -0.25$, and check whether there is a probable technological driver behind the decline. For this we first consult the 1986-87 OOH, and if we find nothing there, we check in the 1996 OOH (BLS 1996), and if we

61. We have excluded four OOH occupations that were merged with or split into an unknown number of occupations: "71 Electroencephalographic technologists and technicians", "203 Public administration — chief executives, legislators, and general administrators", "226 Customer service representatives, utilities" and "293 Electric meter installers and repairers".

still find nothing, we check the 2006 version (BLS 2006).[62] Each OOH occupation thus is assigned an indicator variable for technological-related decline, which equals zero whenever $\hat{g}_k \geq -0.25$, and may equal zero or one when $\hat{g}_k < -0.25$.

We can then decompose the employment growth assigned to each Swedish YRKE5 occupation as follows:

$$\hat{g}_s \equiv \underset{1\times K}{\alpha_s} \times \underset{K\times K}{\mathbb{1}\{\text{technology}\}} \times \underset{K\times 1}{\hat{g}} \quad + \quad \underset{1\times K}{\alpha_s} \times \big(I - \underset{K\times K}{\mathbb{1}\{\text{technology}\}}\big) \times \underset{K\times 1}{\hat{g}}, \tag{3.9}$$

where $\mathbb{1}\{\text{technology}\}$ is a diagonal matrix with the indicator for technologically-related decline on the diagonal, and I is the identity matrix. We define a Swedish YRKE5 as having undergone technology-related decline if it is classified as declining and if the first component of the decomposition (3.9) is less than -0.25, formally

$$[\text{Declining (technology)}]_s \equiv \mathbb{1}\{\alpha_s \times \hat{g} < -0.25 \text{ and } \alpha_s \times \mathbb{1}\{\text{technology}\} \times \hat{g} < -0.25\}.$$

Classification of replaceable occupations To identify which Swedish 5-digit occupations were subject to replacement by specific technologies we separately use a manual approach and an algorithmic approach. In the manual procedure, we asked ourselves whether we could think of a technology that replaces nearly all tasks for each occupation. If yes, the value is one and we write down the technology responsible. If no, the value is zero.

In the algorithmic procedure, we use a pre-defined search procedure on Google. For each occupation, we translate the occupation's title from Swedish to English. Within each YRKE3 category, there is an YRKE5 occupation that encompasses workers within the 3-digit occupation that are not covered by any of the YRKE5 occupations (e.g. "Other within 103").[63] These occupations were also excluded from this procedure, and we return to them below.

For each of the 847 remaining 5-digit occupations, we conduct a Google search to identify articles that pinpointed a specific technology. The search query comprised of three parts and excluded results from replacedbyrobot.info, which considers only anticipated future changes. The query parts were: (i) the exact phrase of the job title, with any necessary modifications as detailed below; (ii) the exact

62. There were four heavily declining ($\hat{g}_k < -0.25$) OOH occupations where we found no information in the OOH editions of 1986, 1996, or 2006, but we still suspected technologically-related decline. Therefore, we searched in other editions of the OOH and other sources, and found potential technological drivers of occupational decline:

213 Radio operators	"Laboursaving [sic] technical advances such as computer-controlled programming and remotely controlled transmitters" (regarding Broadcast and sound engineering technicians and radio operators, BLS 2004:260)
254 Telegraph and teletype operators	Automatic routing of calls, voice message systems (regarding Telephone operators, BLS 1994:291)
346 Motion picture projectionists	Digital projection (Hess 2014)
391 Service station attendants	Self-service pumps at petrol stations (Emek Basker and Klimek 2015)

63. For some YRKE3 classifications, these workers were split into two "other" categories, in this case the later procedures is implemented for each sub-category separately. Furthermore, some YRKE3 classifications are themselves workers "not otherwise classified" within an YRKE2 occupation; here we use the employment-weighted average of all YRKE5 occupations within the same YRKE2.

phrase "replaced by"; (iii) and the word "jobs". The first page of Google search results (10 results) as of December 2019 were consulted to whether an occupation had been replaced; future looking articles, job postings, and forum posts were not considered.

While in most cases, the first part of the query was the exact translated job title, some modifications were necessary. For job titles which included parentheticals (for example "Cleaner (public spaces, offices etc.)"), the parentheticals were excluded from the search terms (i.e. only "Cleaner" was included in the query). Others included an 'or' (for example "Mail sorting clerk or postman"), which were used to separate the job title into two queries (i.e. "Mail sorting clerk" and "postman" were searched separately). This procedure results in an indicator for replacement for each 5-digit occupations that we searched.

For an occupation to be considered as replaced by technology, the following criteria must be met: a specific technology must be identified as materially affecting all of a worker's job duties, and the technology must be in commercial use during the period we study.

Finally, we return to the 206 above-mentioned 5-digit occupations representing miscellaneous workers within each YRKE3 occupation (which together account for 16% of total employment in 1985). For each of these 5-digit occupations we calculate the replacement index as the weighted mean level of replacement in the 3-digit Swedish occupation that contains them, as calculated above, where the weights are the employment-weighted shares of 5-digit occupations that we have classified above within each 3-digit occupations. We implement this procedure for both the algorithmic and manual versions of replacement.

Swedish micro-level variables

In addition to the occupational data, we construct several variables that characterise workers' career outcomes spanning the years 1986-2013; that is, starting with the first year after we measure treatment and ending with the last year available in our data. We start by simply summing up years observed as employed and real annual labour earnings, obtaining the variables cumulative years employed and cumulative earnings.[64] Following Autor, Dorn, Hanson, and Song (2014), we also create a normalised measure of cumulative earnings, whereby we divide cumulative earnings by predicted initial earnings. Cumulative earnings normalised in this way thus give the multiple of (predicted) initial earnings that a worker receives during 1986-2013.[65] We consider further earnings measures—such as rank, discounted cumulative earnings, and earnings growth—in robustness checks.

We construct predicted cumulative earnings based on occupation-specific life-cycle profiles as follows. Separately for each 3-digit occupation, we regress log income on a quartic in age and dummies for sex, county, and education. We generate predicted log income growth in each year by summing

64. We define a worker as employed in a given year if they are identified as working in November (when employment status is measured for the purposes of LISA) of that year and if annual earnings during that year are no lower than the base amount. When we do not observe an individual in a given year—due to emigration or death—we set employment and earnings to zero.

65. The prediction comes from a regression of log earnings on a quartic in age and dummies for gender, county, and seven education categories, run separately for each 3-digit SSYK96 occupation in 1985. We divide by predicted rather than actual initial earnings to eliminate transitory earnings variation, which would introduce an important role for mean reversion into the distribution of normalised cumulative earnings. Autor, Dorn, Hanson, and Song (2014) divide cumulative earnings by earnings averaged across four pre-treatment years for the same reason. Since we do not have annual earnings information prior to 1985, we normalise by predicted earnings instead.

mean real wage growth and the expected growth based on aging, calculated by applying the occupation-specific coefficients on the quartic in age from the previous step. We calculate predicted income in each year by adding predicted log income growth to the predicted log base year income and exponentiating. Cumulative predicted income is the sum of predicted income from 1986–2013. This is the variable we add to the regression when we say we control for life-cycle profiles.

Our measure of long-run occupational mobility is a dummy variable equaling one if the individual worked in the same 3-digit SSYK96 occupation in 2013 as 1985. It equals zero if the individual works in a different occupation or is not employed.[66] Using the PES data, we calculate cumulative days spent unemployed and cumulative days spent in retraining during 1992-2013. We define dummy variables for ever unemployed and ever having participated in retraining. As the PES data are not available for 1986-1991, we cannot capture any unemployment or retraining in these early years of our sample period. Finally, we calculate the retirement age, where we define retirement as a continuous spell of zero annual earnings up to and including age 64.[67]

Subsection 3.B.3 Sample Restrictions

Our starting sample contains all individuals born between 1921-1969—hence aged 16-64 (at some point) in 1985—who were employed in November 1985, whose annual earnings in 1985 were no less than the "base amount" (Swedish: *basbelopp*) specified by the social security administration, and about whom we have complete demographic (including education) and labour market information (including industry and occupation). The base amount is used as an accounting unit when calculating benefits, and it is typically equal to about three months' worth of full-time work at the median wage. As we do not observe hours worked or fulltime status, we use the base amount to exclude individuals with little labour market attachment. There are 3,061,051 individuals fulfilling the above criteria.[68] Our *baseline sample* further restricts birth year to 1949-1960, or ages 25-36 in 1985. We drop younger workers as these are less likely to be attached to the labour market and may not yet have settled on an occupation. And we drop middle-aged and older workers from our baseline sample because we want to focus on workers who did not reach retirement age by 2013, the end of our period of study, in our main analysis. We will analyse workers born before 1949 separately.

66. Our measure of occupational mobility does not capture any temporary exits during the intervening years if workers returned to their initial occupation. A limitation of our data is that they are not conducive to studying high-frequency occupational mobility: During the years 1986-1989 and 1991-1995, we do not observe workers' occupation. And during 1996-2004, the SSYK96 variable contains substantially fewer distinct codes than from 2005 onwards.

67. The LISA database includes individuals older than 64 only during later years. As we do not consistently observe individuals beyond age 64, we assume for all years that individuals aged 65 or older have retired.

68. There were 5,281,382 individuals aged 16-64 in Sweden in 1985. Of those, 4,186,512 were employed in November 1985, and among them, 3,648,034 earned no less than the base amount during 1985. The reduction to 3,061,051 is due to missing education, industry, or occupation information, including cases where YRKE5 occupations do not have matches in the OOH.

Appendix 3.C Appendix figures and tables

Subsection 3.C.1 Figures

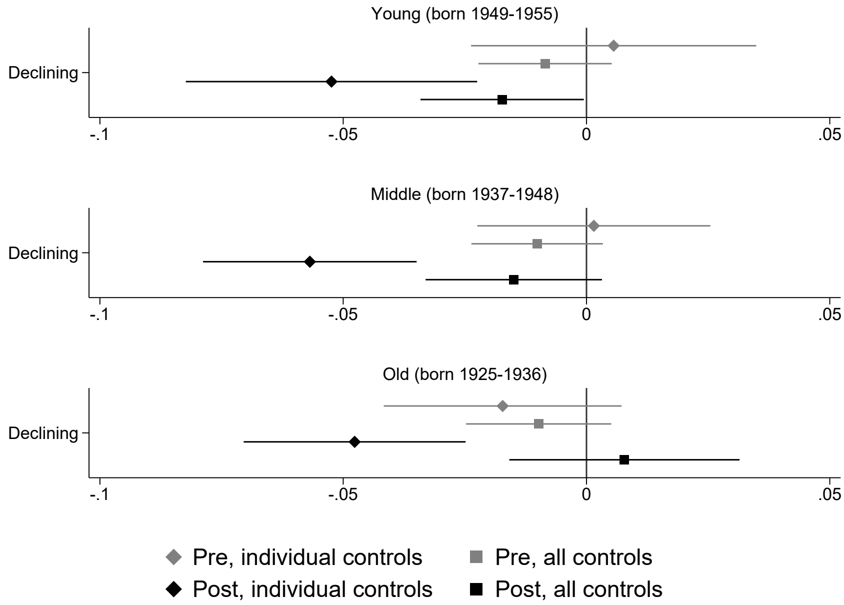

Notes: Coefficients on the declining indicator along with their 95-percent confidence intervals (robust to clustering by 1985 3-digit occupation) are displayed, where the regressions vary the sample, controls, and outcome variables. Coefficients are scaled by the mean of the outcome variable in each estimation sample. 'Post' refers to cumulative earnings 1986-2013. 'Pre' refers to the sum of earnings 1975 & 1980 for the middle and old, and earnings in 1980 for the young. We dropped the 1956-1960 birth cohorts as they did not reach age 25 by 1980, and for a similar reason we did not use 1975 earnings data for the young. 'Individual controls' are those used in column (2) of Table 3.4, and 'all controls' are the ones from column (6) in that table.

Figure 3.4: Earnings prior to occupational decline

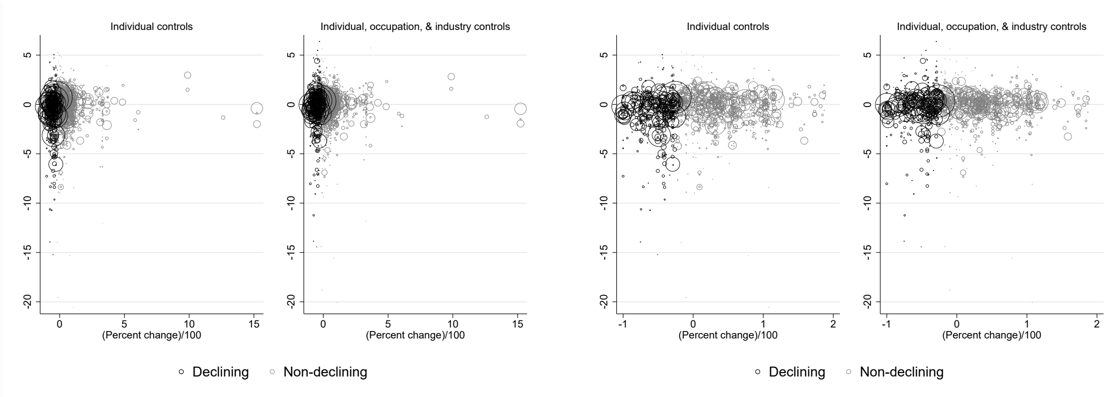

(a) Cumulative employment

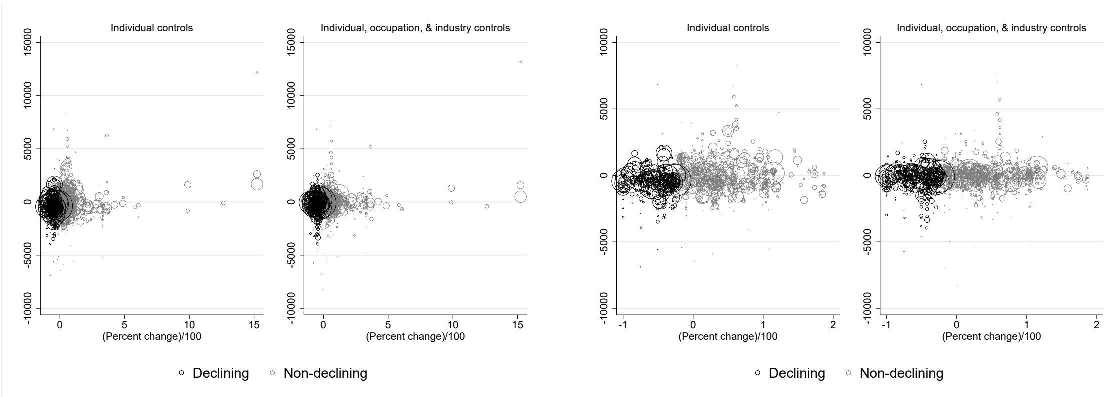

(b) Cumulative earnings

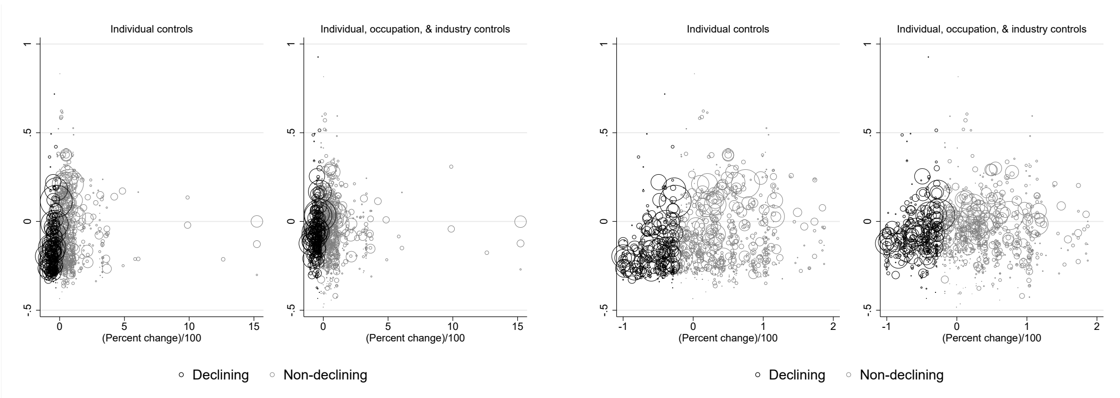

(c) Probability of remaining in the initial 3-digit occupation

Notes: Each bubble represents one of 1,052 5-digit Swedish occupations. Bubbles are scaled according to 1985 Swedish employment. The percent change in employment is assigned based on the changes 1984-2016 in the corresponding US occupations(s). Declining occupations are those that declined by more than 25 percent. Prior to aggregation, outcome variables were residualized based on the regression models in columns (2) and (6) in Tables 3.4 and 3.5, but with 'Declining' times its coefficient added (the mean difference between declining and non-declining occupations in the plots is thus exactly equal to the coefficients reported in the tables). The pairs of graphs on the right are truncated versions of those on the left.

Figure 3.5: Main outcomes and percent change in employment (US)

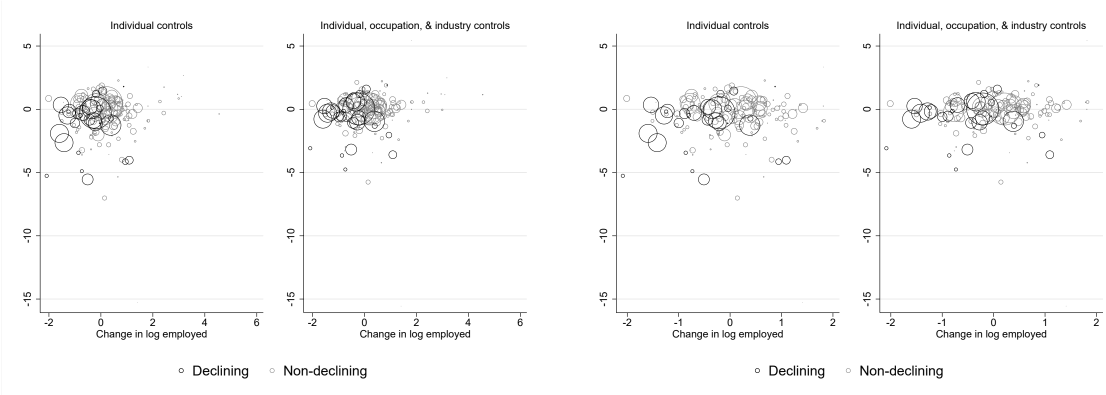

(a) Cumulative employment

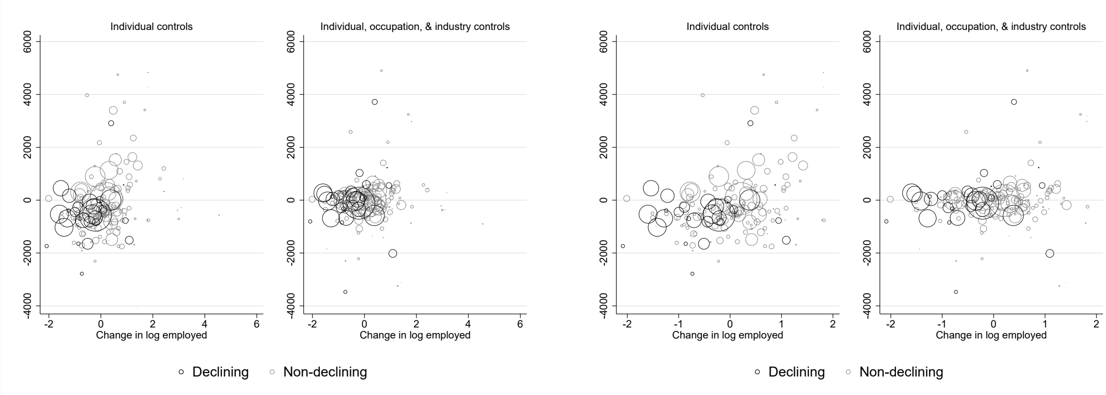

(b) Cumulative earnings

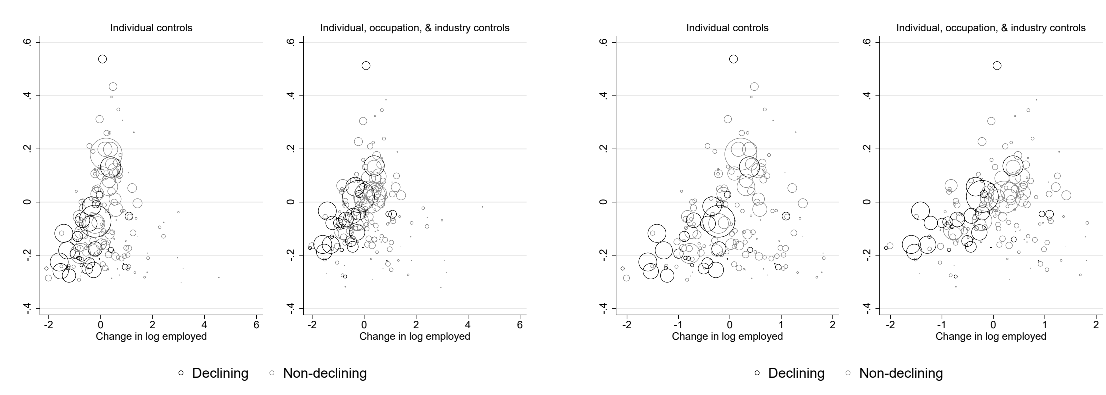

(c) Probability of remaining in the initial 3-digit occupation

Notes: Each bubble represents one of 172 3-digit Swedish occupations. Bubbles are scaled according to 1985 Swedish employment. 'Change in log employment' refers to the actual change in log employment in each Swedish 3-digit occupation from 1985-2013. Occupations marked as declining are those in which more than two thirds of employment in 1985 was in a 5-digit occupation with the 'Declining' indicator equal to one. Prior to aggregation, outcome variables were residualized based on the corresponding regression models reported on in the last panel of Table 3.11, with log employment change times its coefficient added (lines fitted to the plots would thus have slopes equal to the coefficients on log employment change reported in Table 3.11). The pairs of graphs on the right are truncated versions of those on the left.

Figure 3.6: Main outcomes and change in log employment (Sweden)

Subsection 3.C.2 Tables

Further descriptive statistics

Table 3.8: Demographic characteristics across samples

		Ages 16-64			Ages 25-36	
	(1)	(2)	(3)	(4)	(5)	(6)
Female	0.47	0.48	0.49	0.46	0.47	0.48
	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)
Age	39.0	39.3	39.4	30.7	30.8	30.8
	(12.3)	(12.3)	(12.2)	(3.46)	(3.47)	(3.46)
Earnings	190.2	191.2	191.2	181.7	183.0	184.3
	(102.6)	(103.1)	(99.6)	(78.6)	(79.1)	(77.6)
Compulsory school		0.35	0.35		0.25	0.25
		(0.48)	(0.48)		(0.43)	(0.43)
College		0.11	0.096		0.12	0.12
		(0.31)	(0.29)		(0.33)	(0.32)
Observations	3,648,034	3,496,402	3,061,051	1,070,967	995,416	877,324
Dropped if education missing		✓	✓		✓	✓
Dropped if education, occupation, or industry missing			✓			✓

Notes: Samples for columns (1) and (4) include all individuals of the indicated ages residing in Sweden in 1985, who were employed and earned at least one base amount. Samples for the remaining columns are further restricted as indicated by the two bottom rows.

Table 3.9: Employment growth in Swedish 3-digit occupations 1985-2013

	(1)	(2)	(3)	(4)	(5)	(6)
Declining	-0.76				-0.44	-0.46
	(0.17)				(0.18)	(0.18)
Employment share 1985		-1.23			-2.40	-2.31
		(1.61)			(1.57)	(1.53)
Employment growth 1960-1985		0.34			0.16	0.15
		(0.08)			(0.09)	(0.08)
Predicted growth index			0.31		0.22	
			(0.07)		(0.08)	
Prediction: no change				-0.05		0.09
				(0.44)		(0.42)
Prediction: increase, slow				0.46		0.25
				(0.36)		(0.31)
Prediction: increase, average				0.74		0.55
				(0.29)		(0.25)
Prediction: increase, fast				1.13		0.82
				(0.29)		(0.28)
R^2	0.12	0.15	0.21	0.22	0.29	0.29

Notes: The dependent variable is the difference in log employment in Swedish 3-digit occupations between 2013 and 1985. 'Declining' is a binary variable at the level of 1985 Swedish 5-digit occupations indicating employment losses of more than 25 percent over the following three decades in the corresponding US occupation(s). The indicator has been collapsed to the 3-digit level and is thus a continuous regressor. The decline indicator and predictions have been constructed using the Occupational Outlook Handbook (various years). Regressions are weighted by 1985 Swedish employment shares. The number of observations is 172. Robust standard errors in parentheses.

Table 3.10: Alternative cutoffs for occupational decline

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Percent change $\in [-100, -50)$	-0.34	-0.18	-248.1	-90.0	-2.44	-0.98	-0.18	-0.10
	(0.20)	(0.15)	(115.6)	(75.7)	(0.62)	(0.43)	(0.040)	(0.020)
Percent change $\in [-100, -25)$ (baseline)	-0.49	-0.19	-346.6	-126.4	-2.10	-1.11	-0.11	-0.045
	(0.20)	(0.14)	(120.3)	(58.3)	(0.53)	(0.36)	(0.041)	(0.020)
Percent change $\in [-100, 0)$	-0.043	-0.0030	-35.0	-57.5	-0.70	-0.91	-0.15	-0.063
	(0.20)	(0.13)	(158.8)	(74.7)	(0.70)	(0.47)	(0.041)	(0.021)
Percent change $\in [-100, 31)$ (below median)	0.14	0.15	-46.5	-61.9	-0.55	-0.53	-0.087	-0.0094
	(0.18)	(0.13)	(150.7)	(76.1)	(0.57)	(0.50)	(0.037)	(0.022)
Baseline; control: percent change $\in (-25, 31)$	-0.72	-0.27	-460.5	-126.6	-2.40	-1.17	-0.077	-0.054
	(0.22)	(0.16)	(123.3)	(61.9)	(0.51)	(0.40)	(0.038)	(0.018)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Observations			877,324				553,787	

Notes: Results from regressions of various outcomes on indicators for occupational employment changes to lie in the indicated ranges are shown. Each panel represents a separate set of regressions. The underlying variable is the percentage change in employment for the US occupation(s) corresponding to the Swedish 5-digit occupation that the individual worked in during 1985. The last panel only keeps observations with a percentage change below the median, and the number of observations is thus halved. normalised earnings are cumulative earnings divided by initial predicted earnings. See the notes to Tables 3.4 and 3.5 for further descriptions of variables and sample definitions. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.11: Using continuous occupational employment changes as regressors

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Percent employment change / 100 (US)	-0.019	-0.026	103.7	64.7	0.47	0.25	0.0058	-0.0020
	(0.037)	(0.036)	(30.2)	(14.9)	(0.11)	(0.13)	(0.0068)	(0.0029)
Percent employment change / 100 (US), winsorized	0.010	0.000016	83.8	91.1	0.86	0.46	0.050	0.0035
	(0.11)	(0.080)	(112.0)	(47.5)	(0.40)	(0.25)	(0.025)	(0.014)
Log employment change (SWE)	-0.034	0.049	306.5	73.6	0.85	0.086	0.11	0.066
	(0.15)	(0.11)	(135.1)	(65.9)	(0.50)	(0.50)	(0.031)	(0.017)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Observations	877,324						553,787	

Notes: Results from regressions of various outcomes on change in occupational employment are shown. Each panel represents a separate set of regressions. 'Percent employment change (US)' refers to the percentage change in employment 1984-2016 for the US occupation(s) corresponding to the Swedish 5-digit occupation that the individual worked in during 1985. The winsorized measure of this variable top-codes changes at plus 217 percent (the 95th percentile). 'Log employment change (SWE)' refers to the change in log number employed 1985-2013 in the Swedish 3-digit occupation that the individual works in during 1985. normalised earnings are cumulative earnings divided by initial predicted earnings. See the notes to Tables 3.4 and 3.5 for further descriptions of variables and sample definitions. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.12: Controlling for starting firm fixed effects

	(1)	(2)	(3)	(4)
A. Exposure to occupational employment growth (mean: 0.03)				
Declining	-0.40	-0.25	-0.22	-0.19
	(0.11)	(0.11)	(0.10)	(0.11)
B. Cumulative employment (mean: 23.4)				
Declining	-0.49	-0.37	-0.19	-0.28
	(0.20)	(0.14)	(0.14)	(0.11)
C. Cumulative earnings (mean: 6,926)				
Declining	-347	-476	-126	-169
	(120)	(104)	(58)	(43)
D. Cumulative earnings, normalized (mean: 38.7)				
Declining	-2.10	-2.18	-1.11	-1.45
	(0.53)	(0.40)	(0.36)	(0.34)
E. Remain in starting occupation (mean: 0.29)				
Declining	-0.11	-0.039	-0.045	-0.031
	(0.041)	(0.018)	(0.020)	(0.016)
Individual controls	✓	✓	✓	✓
Occupation-industry controls			✓	✓
Firm fixed effects		✓		✓

Notes: Results from regressions of the indicated outcomes on the indicator for occupational decline are shown. The samples are the same as those used in Tables 3.4 and 3.5. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. Columns (2) and (4) additionally control for starting (1985) firm fixed effects.Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.13: Alternative functional forms for earnings

A. Discounted cumulative earnings

	Discounted cumulative earnings				Discounted cumulative earnings, normalized			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Declining	-152.7	-47.8	-49.5	-33.2	-0.94	-0.47	-0.51	-0.40
	(57.1)	(25.5)	(24.4)	(29.5)	(0.25)	(0.16)	(0.16)	(0.18)
Declining × rank			213.9				1.22	
			(68.5)				(0.25)	
Declining × bottom tercile				-166.5				-0.96
				(47.6)				(0.23)
Declining × top tercile				109.3				0.64
				(62.0)				(0.21)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓	✓	✓		✓	✓	✓
Mean of dep. var.			3,476				19.4	
Mean of dep. var., bottom			2,954				17.5	

B. Rank, logs, and growth

	Percentile rank in cumulative earnings				Logarithm of cumulative earnings				Percent growth in earnings 1985-2013			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Declining	-1.47	-0.85	-0.85	-0.93	-0.060	-0.021	-0.026	-0.00054	-11.7	-4.88	-4.55	1.34
	(0.84)	(0.54)	(0.51)	(0.62)	(0.022)	(0.013)	(0.014)	(0.017)	(3.34)	(2.18)	(2.02)	(2.39)
Declining × rank			5.33				0.17				39.8	
			(0.99)				(0.035)				(8.25)	
Declining × bottom tercile				-3.41				-0.15				-32.6
				(0.80)				(0.037)				(7.44)
Declining × top tercile				3.48				0.072				11.7
				(0.90)				(0.018)				(4.92)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓	✓	✓		✓	✓	✓		✓	✓	✓
Mean of dep. var.			50.5				8.6				69	
Mean of dep. var., bottom			43.0				8.4				127	

Notes: Results from regressions of the indicated earnings measures on the declining indicator, within-occupation earnings rank or tercile dummies (coefficients omitted from table), and their interactions are shown. All regressions control for the level of 1985 earnings, with the exception of rank and logarithm as the outcome variables, in which case 1985 earnings rank and log of 1985 earnings are controlled for, respectively. Discounted cumulative earnings are calculated using an interest rate of 5 percent. normalised earnings are cumulative earnings divided by initial predicted earnings. See the notes to Tables 3.4 and 3.5 for further descriptions of variables and sample definitions. The number of observations is 877,324, except when the log of cumulative earnings is the outcome variable, in which case the number is 875,830. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.14: Heterogeneity by within-occupation residualized earnings rank

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Linear interaction								
Declining	-0.59	-0.20	-332.3	-154.0	-2.32	-1.22	-0.11	-0.042
	(0.22)	(0.14)	(90.2)	(59.3)	(0.56)	(0.37)	(0.041)	(0.020)
Declining × rank	0.92	0.96	407.7	439.5	2.33	2.41	-0.020	-0.014
	(0.33)	(0.29)	(141.9)	(137.3)	(0.59)	(0.56)	(0.016)	(0.015)
B. Dummy interactions								
Declining	-0.26	0.048	-302.5	-94.5	-1.94	-0.92	-0.095	-0.032
	(0.22)	(0.16)	(96.4)	(62.9)	(0.52)	(0.38)	(0.050)	(0.025)
Declining × bottom tercile	-1.16	-1.11	-370.4	-390.7	-2.14	-2.10	-0.015	-0.0084
	(0.36)	(0.33)	(93.5)	(86.2)	(0.48)	(0.44)	(0.019)	(0.017)
Declining × top tercile	0.16	0.24	220.2	202.7	0.99	1.12	-0.037	-0.026
	(0.15)	(0.15)	(109.2)	(111.3)	(0.47)	(0.45)	(0.027)	(0.020)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Mean of dep. var.	23.4		6,926		38.7		0.29	
Mean of dep. var., bottom	22.0		6,139		34.3		0.26	
Observations			877,324				553,787	

Notes: The notes to Table 3.6 apply, with the only difference that rank and terciles refer to the within-occupation distribution of 1985 earnings residualized by gender, cohort, and county. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.15: Heterogeneity by overall earnings rank

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Linear interaction								
Declining	-0.67	-0.29	-347.4	-153.9	-2.36	-1.26	-0.12	-0.047
	(0.20)	(0.13)	(116.4)	(64.7)	(0.59)	(0.39)	(0.042)	(0.019)
Declining × rank	1.55	1.33	131.5	294.9	2.41	1.83	0.068	0.023
	(0.38)	(0.33)	(196.4)	(154.0)	(0.82)	(0.71)	(0.035)	(0.019)
B. Dummy interactions								
Declining	-0.19	0.18	-78.6	6.25	-1.68	-0.46	-0.10	-0.044
	(0.19)	(0.16)	(86.4)	(60.3)	(0.57)	(0.38)	(0.050)	(0.022)
Declining × bottom tercile	-2.02	-1.83	-473.9	-425.0	-2.61	-2.51	-0.080	-0.020
	(0.51)	(0.43)	(139.8)	(132.8)	(0.79)	(0.73)	(0.037)	(0.016)
Declining × top tercile	0.30	0.14	-374.2	-110.8	0.39	-0.26	0.018	0.0092
	(0.16)	(0.16)	(143.7)	(87.3)	(0.58)	(0.43)	(0.042)	(0.024)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Mean of dep. var.	23.4		6,926		38.7		0.29	
Mean of dep. var., bottom	21.9		5,058		35.2		0.29	
Observations	877,324						553,787	

Notes: The notes to Table 3.6 apply, with the only difference that rank and terciles refer to the overall distribution of 1985 earnings. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.16: Heterogeneity by overall earnings rank—residualized earnings rank

	Employment		Earnings		Earnings, normalized		Remain	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
A. Linear interaction								
Declining	-0.53	-0.20	-346.5	-135.5	-2.16	-1.14	-0.11	-0.044
	(0.20)	(0.14)	(118.8)	(62.2)	(0.55)	(0.37)	(0.041)	(0.020)
Declining × rank	0.99	0.91	211.4	295.0	1.36	1.42	-0.017	-0.016
	(0.36)	(0.30)	(186.3)	(142.6)	(0.76)	(0.67)	(0.028)	(0.020)
B. Dummy interactions								
Declining	-0.16	0.13	-224.3	-51.2	-1.76	-0.64	-0.11	-0.045
	(0.17)	(0.15)	(102.9)	(67.5)	(0.52)	(0.38)	(0.050)	(0.026)
Declining × bottom tercile	-1.15	-1.06	-391.2	-352.9	-1.48	-1.76	0.011	0.013
	(0.37)	(0.30)	(105.2)	(112.7)	(0.59)	(0.58)	(0.020)	(0.022)
Declining × top tercile	-0.027	-0.054	-29.0	62.7	0.19	0.041	-0.011	-0.0089
	(0.15)	(0.16)	(162.2)	(108.7)	(0.58)	(0.47)	(0.037)	(0.025)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓
Mean of dep. var.	23.4		6,926		38.7		0.29	
Mean of dep. var., bottom	21.3		5,295		32.3		0.24	
Observations			877,324				553,787	

Notes: The notes to Table 3.6 apply, with the only difference that rank and terciles refer to the overall distribution of 1985 residualized earnings (residualized by gender, cohort, and county). Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.17: Cumulative earnings of leavers and stayers in declining and non-declining occupations

	(1)	(2)	(3)	(4)
A. All workers (553,170 observations)				
Remain	335	303	305	284
	(122)	(91)	(133)	(101)
Declining			-271	-126
			(122)	(90)
Declining × remain			177	189
			(239)	(184)
B. Employed in 2013 (404,046 observations)				
Remain	-398	-498	-439	-531
	(115)	(66)	(124)	(72)
Declining			-356	-187
			(123)	(94)
Declining × remain			238	311
			(231)	(158)
C. Employed in 2013, bottom third (140,893 observations)				
Remain	-109	-286	-133	-307
	(139)	(85)	(143)	(81)
Declining			-417	-237
			(145)	(173)
Declining × remain			-33	233
			(596)	(425)
Individual controls	✓	✓	✓	✓
Occupation & industry controls		✓		✓

Notes: The dependent variable is cumulative earnings 1986-2013 in thousands of 2014 SEK. 'Remain' is an indicator for working in the same 3-digit occupation in 2013 as in 1985. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. The sample is the same as that in Table 3.5, except for the restrictions indicated. Sampling weights are applied. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.18: Occupational decline and individual-level cumulative employment and earnings 1986-2013—'doughnut' specifications

	(1)	(2)	(3)	(4)	(5)	(6)
A. Cumulative years employed 1986-2013 (mean: 23.5)						
Declining	-1.46	-0.97	-0.97	-0.82	-0.35	-0.41
	(0.53)	(0.42)	(0.42)	(0.46)	(0.28)	(0.29)
B. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,612)						
Declining	-484	-403	-333	-140	-81	-217
	(608)	(196)	(177)	(181)	(158)	(167)
C. Cumulative real earnings divided by predicted initial earnings (mean: 39.2)						
Declining	-5.40	-2.49	-2.56	-1.81	-1.18	-1.69
	(1.33)	(1.09)	(0.98)	(1.07)	(0.82)	(1.05)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1985 in a subsequently declining occupation are shown. The sample is the same as in Table 3.4, but excludes 3-digit occupations in which some but not all 5-digit occupations are coded as declining. Thus, within each 3-digit occupation, either all 5-digit sub-occupations decline, or none, leaving out intermediate cases ('doughnut'). The number of observations is 488,484. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.19: Occupational decline and individual occupational stability—'doughnut' specifications

	(1)	(2)	(3)	(4)	(5)	(6)
A. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.35)						
Declining	-0.25	-0.21	-0.21	-0.12	-0.17	-0.10
	(0.046)	(0.051)	(0.052)	(0.044)	(0.046)	(0.046)
B. Probability of working in same 2-digit occupation in 2013 as in 1985 (mean: 0.40)						
Declining	-0.21	-0.16	-0.16	-0.090	-0.12	-0.059
	(0.038)	(0.045)	(0.046)	(0.043)	(0.045)	(0.042)
C. Probability of working in same 1-digit occupation in 2013 as in 1985 (mean: 0.44)						
Declining	-0.19	-0.14	-0.14	-0.077	-0.11	-0.045
	(0.036)	(0.042)	(0.043)	(0.042)	(0.043)	(0.033)
Demographics & earnings		✓	✓	✓	✓	✓
Life-cycle profiles			✓	✓	✓	✓
Predictors of growth				✓	✓	✓
Occupation dummies					✓	✓
Industry dummies						✓

Notes: Results from regressions of the indicated outcomes on a dummy for working in 1985 in a subsequently declining occupation are shown. The sample is the same as in Table 3.5, but with the 'doughnut' restrictions from Table 3.18 applied. The number of observations is 333,357. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.20: Occupational decline and gender

	Employment				Earnings				Remaining			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A1. Women, linear interaction												
Declining	-0.98	-0.21	-0.94	-0.22	-66.5	-15.2	-297.0	33.8	-0.22	-0.061	-0.22	-0.066
	(0.26)	(0.17)	(0.29)	(0.18)	(144.1)	(79.0)	(140.0)	(97.6)	(0.072)	(0.023)	(0.065)	(0.024)
Declining × rank			0.021	0.017			96.8	179.1			-0.048	-0.013
			(0.34)	(0.27)			(129.6)	(162.7)			(0.021)	(0.014)
A2. Women, dummy interactions												
Declining	-0.98	-0.21	-1.05	-0.34	-66.5	-15.2	-274.6	-61.6	-0.22	-0.061	-0.23	-0.077
	(0.26)	(0.17)	(0.29)	(0.24)	(144.1)	(79.0)	(123.6)	(71.8)	(0.072)	(0.023)	(0.070)	(0.026)
Declining × bottom tercile			0.17	0.23			32.9	25.9			0.040	0.024
			(0.26)	(0.23)			(143.3)	(136.4)			(0.022)	(0.020)
Declining × top tercile			-0.011	0.057			121.8	201.5			-0.019	0.017
			(0.43)	(0.37)			(193.7)	(212.3)			(0.020)	(0.018)
Mean of dep. var.		24.1				5,921				0.32		
Mean of dep. var., bottom		23.7				5,684				0.29		
B1. Men, linear interaction												
Declining	-0.18	0.0055	-0.35	-0.13	-261.7	-110.6	-253.9	-142.7	-0.067	-0.029	-0.061	-0.026
	(0.23)	(0.14)	(0.28)	(0.16)	(116.6)	(55.7)	(107.1)	(54.4)	(0.038)	(0.020)	(0.039)	(0.020)
Declining × rank			0.73	0.60			299.1	148.7			-0.038	-0.027
			(0.40)	(0.30)			(146.8)	(116.8)			(0.034)	(0.023)
B2. Men, dummy interactions												
Declining	-0.18	0.0055	-0.26	-0.0062	-261.7	-110.6	-232.0	-79.3	-0.067	-0.029	-0.029	0.0015
	(0.23)	(0.14)	(0.26)	(0.16)	(116.6)	(55.7)	(123.0)	(58.1)	(0.038)	(0.020)	(0.042)	(0.022)
Declining × bottom tercile			-0.67	-0.70			-259.7	-218.2			-0.040	-0.036
			(0.40)	(0.37)			(113.3)	(82.6)			(0.022)	(0.022)
Declining × top tercile			0.45	0.38			143.8	23.3			-0.070	-0.055
			(0.24)	(0.18)			(114.3)	(95.1)			(0.033)	(0.022)
Mean of dep. var.		23.5				8,042				0.25		
Mean of dep. var., bottom		21.5				6,869				0.24		
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓		✓		✓

Notes: The notes to Table 3.6 apply. Results are from regressions run separately by gender. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.21: Occupational decline and gender, residualized earnings rank

	Employment				Earnings				Remaining			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A1. Women, linear interaction												
Declining	-0.98	-0.21	-1.04	-0.26	-66.5	-15.2	-165.7	-63.4	-0.22	-0.061	-0.22	-0.057
	(0.26)	(0.17)	(0.25)	(0.18)	(144.1)	(79.0)	(106.0)	(73.2)	(0.072)	(0.023)	(0.070)	(0.023)
Declining × rank			0.37	0.56			220.0	301.3			-0.025	-0.024
			(0.28)	(0.26)			(136.3)	(135.4)			(0.017)	(0.015)
A2. Women, dummy interactions												
Declining	-0.98	-0.21	-0.78	-0.078	-66.5	-15.2	-67.9	-29.1	-0.22	-0.061	-0.24	-0.068
	(0.26)	(0.17)	(0.30)	(0.29)	(144.1)	(79.0)	(81.6)	(88.7)	(0.072)	(0.023)	(0.075)	(0.027)
Declining × bottom tercile			-0.65	-0.63			-303.1	-251.4			0.043	0.036
			(0.34)	(0.33)			(160.2)	(128.2)			(0.020)	(0.018)
Declining × top tercile			-0.096	0.083			75.1	174.8			0.0093	-0.0049
			(0.27)	(0.28)			(157.2)	(164.6)			(0.016)	(0.017)
Mean of dep. var.		24.1				5,921				0.32		
Mean of dep. var., bottom		23.6				5,524				0.30		
B1. Men, linear interaction												
Declining	-0.18	0.0055	-0.27	0.073	-261.7	-110.6	-246.7	-85.0	-0.067	-0.029	-0.067	-0.027
	(0.23)	(0.14)	(0.26)	(0.14)	(116.6)	(55.7)	(105.9)	(54.3)	(0.038)	(0.020)	(0.038)	(0.019)
Declining × rank			0.73	0.69			62.2	45.3			-0.0092	-0.0030
			(0.36)	(0.30)			(148.2)	(118.1)			(0.025)	(0.018)
B2. Men, dummy interactions												
Declining	-0.18	0.0055	-0.091	0.16	-261.7	-110.6	-193.3	-23.3	-0.067	-0.029	-0.051	-0.018
	(0.23)	(0.14)	(0.24)	(0.16)	(116.6)	(55.7)	(112.0)	(57.6)	(0.038)	(0.020)	(0.049)	(0.024)
Declining × bottom tercile			-0.85	-0.80			-152.7	-147.5			-0.023	-0.020
			(0.38)	(0.34)			(95.6)	(79.4)			(0.025)	(0.020)
Declining × top tercile			0.28	0.32			-52.6	-75.6			-0.028	-0.015
			(0.20)	(0.20)			(118.9)	(99.8)			(0.039)	(0.028)
Mean of dep. var.		23.5				8,042				0.25		
Mean of dep. var., bottom		21.6				6,858				0.23		
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓		✓		✓

Notes: The notes to Table 3.14 apply. Results are from regressions run separately by gender. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Further outcomes

Table 3.22: Occupational decline and older workers

	Cumulative years employed			Cumulative earnings			Age at retirement		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
A. Workers aged 37-48 in 1985 (976,637 observations)									
Declining	-0.70	-0.32	-0.47	-273.1	-72.9	-99.4	-0.31	-0.13	-0.20
	(0.16)	(0.11)	(0.12)	(53.0)	(43.4)	(39.7)	(0.078)	(0.054)	(0.060)
Declining × rank			0.98			173.7			0.43
			(0.25)			(85.8)			(0.13)
Mean of dependent variable		17.2			4,759			63.1	
B. Workers aged 49-60 in 1985 (650,538 observations)									
Declining	-0.29	-0.047	-0.087	-75.0	12.4	8.09	-0.16	-0.0038	-0.025
	(0.085)	(0.070)	(0.072)	(18.2)	(18.8)	(18.2)	(0.057)	(0.045)	(0.046)
Declining × rank			0.18			14.4			0.091
			(0.093)			(26.4)			(0.059)
Mean of dependent variable		7.0			1,576			63.7	
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓	✓		✓	✓		✓	✓

Notes: Results from regressions of the indicated outcomes on the declining indicator, within-occupation earnings rank (coefficient omitted from table), and their interaction are shown. Retirement is defined as the beginning of a continuous spell of years with zero earnings lasting until age 65. Samples are as in panel A of Table 3.2, but restricted by age as indicated. See the notes to Table 3.6 for a description of right-hand side variables. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.23: Occupational decline and geographic stability

	Municipality				Commuting zone				County			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A. Linear interaction												
Declining	0.38	-0.70	0.32	-0.74	-0.47	-1.07	-0.53	-1.12	-0.23	-0.86	-0.28	-0.91
	(0.84)	(0.63)	(0.80)	(0.62)	(0.75)	(0.61)	(0.74)	(0.60)	(0.69)	(0.55)	(0.68)	(0.55)
Declining × rank			2.80	2.88			2.58	2.73			2.28	2.39
			(1.09)	(1.02)			(0.93)	(0.85)			(0.86)	(0.79)
B. Dummy interactions												
Declining	0.38	-0.70	0.54	-0.52	-0.47	-1.07	-0.55	-1.12	-0.23	-0.86	-0.11	-0.72
	(0.84)	(0.63)	(1.09)	(0.81)	(0.75)	(0.61)	(0.92)	(0.74)	(0.69)	(0.55)	(0.83)	(0.64)
Declining × bottom tercile			-2.39	-2.46			-1.83	-1.97			-1.93	-2.04
			(1.18)	(1.16)			(0.98)	(0.95)			(0.96)	(0.92)
Declining × top tercile			1.73	1.79			1.92	2.00			1.43	1.49
			(0.82)	(0.71)			(0.71)	(0.62)			(0.62)	(0.54)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓		✓		✓
Mean of dep. var.		66				81				82		
Mean of dep. var., bottom		65				79				81		

Notes: Outcome variables are indicators for residing in the same location in 2013 as in 1985. Means, coefficients, and standard errors have been multiplied by 100 for readability. The notes to Table 3.6 apply. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.24: Occupational decline and geographic stability—residualized earnings rank

	Municipality				Commuting zone				County			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
A. Linear interaction												
Declining	0.38	-0.70	0.075	-0.80	-0.47	-1.07	-0.75	-1.15	-0.23	-0.86	-0.48	-0.92
	(0.84)	(0.63)	(0.74)	(0.63)	(0.75)	(0.61)	(0.71)	(0.62)	(0.69)	(0.55)	(0.64)	(0.56)
Declining × rank			2.13	2.18			1.95	2.17			1.77	1.92
			(0.92)	(0.96)			(0.82)	(0.81)			(0.80)	(0.80)
B. Dummy interactions												
Declining	0.38	-0.70	1.00	-0.071	-0.47	-1.07	0.10	-0.49	-0.23	-0.86	0.40	-0.22
	(0.84)	(0.63)	(0.88)	(0.67)	(0.75)	(0.61)	(0.73)	(0.61)	(0.69)	(0.55)	(0.67)	(0.54)
Declining × bottom tercile			-2.77	-2.70			-2.61	-2.63			-2.55	-2.55
			(0.98)	(0.99)			(0.87)	(0.84)			(0.87)	(0.85)
Declining × top tercile			0.10	0.29			0.11	0.41			-0.046	0.19
			(0.54)	(0.57)			(0.50)	(0.54)			(0.45)	(0.49)
Individual controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓		✓		✓		✓		✓
Mean of dep. var.		66				81				82		
Mean of dep. var., bottom		63				78				79		

Notes: Outcome variables are indicators for residing in the same location in 2013 as in 1985. Means, coefficients, and standard errors have been multiplied by 100 for readability. The notes to Table 3.14 apply. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Alternative treatments

Table 3.25: Baseline characteristics of workers in subsequently declining occupations—technology-related declines

	(1) Female	(2) age	(3) Compulsory school	(4) High school	(5) College	(6) Earnings	(7) Manufacturing
A. Occupational decline, pooled							
Intercept	0.52	39.5	0.33	0.56	0.11	191.3	0.25
	(0.078)	(0.41)	(0.030)	(0.033)	(0.027)	(10.8)	(0.050)
Declining	-0.25	-0.89	0.13	-0.063	-0.070	-0.23	0.38
	(0.088)	(0.63)	(0.035)	(0.034)	(0.028)	(11.0)	(0.085)
B. Occupational decline, by presence of technology link							
Intercept	0.52	39.5	0.33	0.56	0.11	191.3	0.25
	(0.078)	(0.41)	(0.030)	(0.033)	(0.027)	(10.8)	(0.050)
Declining	-0.32	0.033	0.13	-0.086	-0.041	5.31	0.26
	(0.10)	(0.87)	(0.056)	(0.051)	(0.035)	(15.0)	(0.10)
Declining (technology)	0.11	-1.49	0.010	0.037	-0.047	-8.90	0.20
	(0.097)	(1.01)	(0.059)	(0.050)	(0.025)	(14.6)	(0.12)

Notes: Results from OLS regressions of various baseline (1985) characteristics on a constant and indicators for working in a declining occupation are shown (see the notes to Table 3.27 for a description of these indicators). Earnings are measured in thousand Swedish crowns inflated to 2014 levels. The sample is the same as in panel A of Table 3.2. The number of observations is 3,061,051. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.26: Occupational decline and individual-level cumulative employment and earnings 1986-2013—technology-related decline

	(1)	(2)	(3)	(4)	(5)
A. Change in log employment 1985-2013 in 3-digit occupation					
Declining	-0.43 (0.16)	-0.25 (0.09)			
Declining (technology)	0.06 (0.17)	0.05 (0.13)	-0.38 (0.12)	-0.21 (0.14)	
B. Cumulative years employed 1986-2013 (mean: 23.4)					
Declining	-0.93 (0.44)	-0.45 (0.24)			
Declining (technology)	0.72 (0.45)	0.42 (0.23)	-0.21 (0.16)	0.01 (0.14)	-0.16 (0.12)
Declining (tech) × rank					1.31 (0.36)
C. Cumulative real earnings ('000 2014 SEK) 1986-2013 (mean: 6,926)					
Declining	-426 (232)	-181 (93)			
Declining (technology)	128 (262)	87 (102)	-303 (131)	-107 (65)	-122 (61)
Declining (tech) × rank					491 (155)
D. Probability of working in same 3-digit occupation in 2013 as in 1985 (mean: 0.29)					
Declining	-0.077 (0.051)	-0.029 (0.022)			
Declining (technology)	-0.058 (0.044)	-0.025 (0.029)	-0.135 (0.043)	-0.053 (0.026)	-0.056 (0.026)
Declining (tech) × rank					0.019 (0.016)
Individual controls	✓	✓	✓	✓	✓
Occupation & industry controls		✓		✓	✓
Observations (population—sample)	877,324—553,787			836,057—532,422	

Notes: Results from regressions of the indicated outcomes on indicators for working in 1985 in a subsequently declining occupation are shown (see the notes to Table 3.27 for a description of these indicators). Columns (1)-(2) are based on the same samples as the results in Tables 3.4 and 3.5. Columns (3)-(5) exclude workers in occupations that are classified as declining without a technology link. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. Sampling weights are used in the regression reported in panel C. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.

Table 3.27: Employment growth in Swedish 3-digit occupations 1985-2013—technology-related declines

	(1)	(2)	(3)	(4)	(5)	(6)
Declining	-0.76	-0.44	-0.92	-0.37		
	(0.17)	(0.18)	(0.27)	(0.27)		
Declining (technology)			0.27	-0.11	-0.69	-0.49
			(0.33)	(0.35)	(0.20)	(0.25)
Employment share 1985		-2.40		-2.41		-2.28
		(1.57)		(1.56)		(1.61)
Employment growth 1960-1985		0.16		0.16		0.16
		(0.09)		(0.09)		(0.09)
Predicted growth index		0.22		0.23		0.22
		(0.08)		(0.08)		(0.09)
R^2	0.12	0.29	0.12	0.29	0.06	0.22
Observations	172	172	172	172	148	148

Notes: The dependent variable is the difference in log employment in Swedish 3-digit occupations between 2013 and 1985. ‘Declining’ is a binary variable at the level of 1985 Swedish 5-digit occupations indicating employment losses of more than 25 percent over the following three decades in the corresponding US occupation(s). ‘Declining (technology)’ indicates that this decline is related to technological replacement. Both indicators have been collapsed to the 3-digit level and are thus continuous regressors. Columns (10) and (11) exclude 3-digit occupations where ‘Declining’ is larger than or equal to 0.5 and ‘Declining (technology)’ is smaller than 0.5. Decline indicators and predictions have been constructed using the Occupational Outlook Handbook (various years). Regressions are weighted by 1985 Swedish employment shares. Robust standard errors in parentheses.

Table 3.28: Technological replacement, exposure to decline, and career outcomes

	(1)	(2)	(3)	(4)
A. Decline indicator				
Replaced (algorithmic coding)	0.27	0.25		
	(0.16)	(0.11)		
Replaced (manual coding)			0.17	0.11
			(0.15)	(0.15)
B. Occupation-level employment growth				
Replaced (algorithmic coding)	-0.55	-0.28		
	(0.20)	(0.14)		
Replaced (manual coding)			-0.31	0.04
			(0.10)	(0.14)
C. Cumulative employment				
Replaced (algorithmic coding)	-0.39	-0.16		
	(0.68)	(0.31)		
Replaced (manual coding)			-0.32	-0.49
			(0.31)	(0.13)
D. Cumulative earnings				
Replaced (algorithmic coding)	-172	-96		
	(228)	(113)		
Replaced (manual coding)			-479	-328
			(132)	(129)
Individual controls	✓	✓	✓	✓
Occupation-industry controls		✓		✓

Notes: Results from regressions of the indicated outcomes on indicators for occupation-level technological replacement are shown. The sample is the same as that used in Table 3.4. Individual-level controls include female, cohort, county, and education dummies, as well as earnings in 1985. Occupation and industry controls include predicted life-time income, predictors of occupational growth, occupation dummies, and industry dummies. Robust standard errors, clustered by 1985 3-digit occupation, in parentheses.