DISSERTATION SERIES 2023:1

Workers and occupations in a changing labour market

The heterogeneous effects of mass layoffs and social safety nets

Yaroslav Yakymovych

IFAU
INSTITUTE FOR EVALUATION OF LABOUR MARKET AND EDUCATION POLICY

The Institute for Evaluation of Labour Market and Education Policy (IFAU) is a research institute under the Swedish Ministry of Employment, situated in Uppsala. IFAU's objective is to promote, support and carry out scientific evaluations. The assignment includes: the effects of labour market and educational policies, studies of the functioning of the labour market and the labour market effects of social insurance policies. IFAU shall also disseminate its results so that they become accessible to different interested parties in Sweden and abroad.

IFAU also provides funding for research projects within its areas of interest. The deadline for applications is October 1 each year. Since the researchers at IFAU are mainly economists, researchers from other disciplines are encouraged to apply for funding.

IFAU is run by a Director-General. The institute has a scientific council, consisting of a chairman, the Director-General and five other members. Among other things, the scientific council proposes a decision for the allocation of research grants. A reference group including representatives for employer organizations and trade unions, as well as the ministries and authorities concerned is also connected to the institute.

Postal address: P O Box 513, 751 20 Uppsala
Visiting address: Kyrkogårdsgatan 6, Uppsala
Phone: +46 18 471 70 70
Fax: +46 18 471 70 71
ifau@ifau.uu.se
www.ifau.se

Dissertation presented at Uppsala University to be publicly examined in Lecture Hall 2, Ekonomikum, Kyrkogårdsgatan 10, Uppsala, Thursday, 6 October 2022 at 15:15 for the degree of Doctor of Philosophy.
Essay II has been published by IFAU as working paper 2022:15 and Swedish report 2023:1

ISSN 1651-4149

Economic Studies 205

Yaroslav Yakymovych

Workers and Occupations in a Changing Labour Market

The Heterogeneous Effects of Mass Layoffs and Social Safety Nets

Department of Economics, Uppsala University

Visiting address: Kyrkogårdsgatan 10, Uppsala, Sweden
Postal address: Box 513, SE-751 20 Uppsala, Sweden
Telephone: +46 18 471 00 00
Telefax: +46 18 471 14 78
Internet: http://www.nek.uu.se/

ECONOMICS AT UPPSALA UNIVERSITY

The Department of Economics at Uppsala University has a long history. The first chair in Economics in the Nordic countries was instituted at Uppsala University in 1741.

The main focus of research at the department has varied over the years but has typically been oriented towards policy-relevant applied economics, including both theoretical and empirical studies. The currently most active areas of research can be grouped into six categories:

* Labour economics
* Public economics
* Macroeconomics
* Microeconometrics
* Environmental economics
* Housing and urban economics

Yaroslav Yakymovych

Workers and Occupations in a Changing Labour Market

The Heterogeneous Effects of Mass Layoffs and Social Safety Nets

UPPSALA
UNIVERSITET

Dissertation presented at Uppsala University to be publicly examined in Lecture Hall 2, Ekonomikum, Kyrkogårdsgatan 10, Uppsala, Thursday, 6 October 2022 at 15:15 for the degree of Doctor of Philosophy. The examination will be conducted in English. Faculty examiner: Prof. Dr. Michael Lechner (University of St.Gallen).

Abstract
Yakymovych, Y. 2022. Workers and Occupations in a Changing Labour Market. The Heterogeneous Effects of Mass Layoffs and Social Safety Nets. *Economic studies* 205. 212 pp. Uppsala: Department of Economics. ISBN 978-91-506-2967-5.

Essay I: Sickness insurance guarantees employees the right to take leave from work when they are sick, but is vulnerable to excessive use. This paper studies which workers react to changes in monitoring by physicians in a large-scale randomised experiment. I use a causal forest to identify heterogeneous effects on the duration of workers' sickness absence spells. Those who are most sensitive to monitoring have a history of extensive sick leave uptake, low socioeconomic status, and male gender. A targeted monitoring policy is estimated to be 40 percent more efficient than a random one.

Essay II: Routine-biased technological change has depressed prospects for workers in exposed occupations, with those displaced in mass layoffs particularly affected. I compare labour market outcomes of displaced routine workers to those of displaced non-routine workers using Swedish microdata. The results show substantial and persistent routine penalties among displaced workers. A possible channel is the loss of occupation- and industry-specific human capital, as routine workers are unable to find jobs similar to those they had before displacement. I do not find evidence that switching to a non-routine occupation reduces routine workers' losses.

Essay III (with Susan Athey, Lisa Simon, Oskar Nordström Skans and Johan Vikström): We study heterogeneity in the impact of job loss in mass layoffs using generalized random forests. We identify the groups of workers who are hit the hardest and document substantial and persistent variation in displacement losses. Worker attributes and semi-aggregate local and industry conditions interact to generate this heterogeneity. Old and less-educated workers lose six times as much as young and highly educated workers. Nevertheless, there is overlap among the losses of these two groups, much of which is related to industries and locations. Working in manufacturing and living in a rural area are strong predictors of severe displacement losses. No simple rule is as effective at identifying vulnerable workers as the more flexible generalized random forest.

Essay IV (with Adrian Adermon, Simon Ek and Georg Graetz): Using a new identification strategy, we jointly estimate growth in occupational wage premia and time-varying occupation-specific lifecycle profiles for Swedish workers in 1996–2013. We document a substantial increase in between-occupation wage inequality due to differential growth in premia. The association of wage premium growth and employment growth is positive, suggesting that premium growth is predominantly driven by demand side factors. Wage growth due to occupation-specific skill acquisition was more dispersed in the early years of the sample period. Our results are robust to allowing for occupation-level changes in returns to cognitive and psycho-social skills.

Keywords: Structural change; Mass layoffs; Sickness absence; Causal forest

Yaroslav Yakymovych, Department of Economics, Box 513, Uppsala University, SE-75120 Uppsala, Sweden.

ISSN 0283-7668
ISBN 978-91-506-2967-5
URN urn:nbn:se:uu:diva-481848 (http://urn.kb.se/resolve?urn=urn:nbn:se:uu:diva-481848)

To my family

Acknowledgments

As I finish writing up this thesis, I cannot help but reflect on how much has changed since I started working on it in 2017. I have of course changed – both as an economist and as a person. Thanks to my wonderful supervisors, co-authors and colleagues, I now know how to write research papers, navigate the jungles of Swedish microdata and code in Stata in a (somewhat) elegant way. My time as a PhD student has also given me the joy of meeting many great friends, who I hope will continue cheering up my life for years to come.

However, the world around me has perhaps changed even more than I have. Things that were once the stuff of dystopian movies or outlandish fantasies have become reality, upending the boundaries of what we imagined possible. At times, it feels like two of the years that went into writing this thesis just evaporated away, spent as they were avoiding other people while the world was ravaged by a pandemic. This spring, Ukraine, where I was born, was attacked on the whim of a crazy regime. Tens of thousands of people have died in a war fought to turn back the course of history and restore an evil empire. Needless to say, it is not easy to write papers on economics when a country you hold dear fights for its survival and your grandparents and relatives live in a city which can be struck by enemy missiles at any time.

In spite of the dark times, I have managed to finish this thesis, which would have been impossible without invaluable help and support from many incredible people. First and foremost, I would like to thank my main supervisor Stefan Eriksson. He has always been there when I had questions regarding research or administrative technicalities. Stefan's detailed comments and suggestions have enabled my papers to become what they are and our weekly meetings always stimulated me to be productive and move towards the next goal.

I would also like to thank my assistant supervisors, Georg Graetz and Adrian Adermon. Georg has been key in helping me become a proficient Stata coder and provided inestimable support throughout the job market process. Adrian has been key to cooperation with IFAU in my projects, organising data access and navigating me through the steps of working paper publication.

My co-authors, Oskar Nordström Skans, Johan Vikström, Simon Ek, Lisa Simon and Susan Athey, have of course greatly contributed to the papers that we have worked on together. I have learnt extremely much from their knowledge of economics, econometrics and machine learning. I would especially like to thank Oskar for introducing me to the project on mass layoff heterogeneity and for his indispensable help while I was on the job market.

My former supervisor, Magnus Gustavsson, who has left academia, got me set on my first project about mass layoffs and routine workers, for which I am grateful. Magdalena Dominguez and Christoph Hedtrich made the job market less daunting by sharing their experiences and giving tips on how to write applications. Per Johansson graciously allowed access to the data required for my job market paper. Michael Böhm and David Strömberg, who were my licentiate and final seminar opponents respectively, provided very useful comments, which have improved the quality of my work. The input given by department faculty at Uppsala Labour Group presentations, end-of-year seminars and mock job market events has been equally valuable. I would also like to thank the teachers of my first-year and second-year courses, who gave me a strong theoretical and econometric foundation to stand on, as well as a detailed knowledge of the labour economics literature. The administrative staff at the department have always been extremely helpful. Per Engström and Oscar Erixson took me on as an additional teaching assistant for their course when I was still a master's student, setting me on the first step of my career at the university.

Of course, the PhD experience would have been miserable without the friends whom I have met along the way. My cohort members, Josefin Videnord, Anna Johansson-Björkman, Raoul van Maarseveen, Davide Gandolfi, Markus Ridder, Alice Hallman and Adrian Poignant, helped me through the struggle of the first year. We have kept going strong even after the blood, sweat and tears of the introductory courses were over. I have thoroughly enjoyed the board game nights, kubb games, burgers & beers and weddings, and have promptly responded to a number of flower watering requests (to the best of my knowledge, only one plant succumbed on my watch).

I would also like to thank all other PhD students at the department, especially Zeynep, Fei, Elin, Edvin, Anton, Maxi, Jan, Davide C, Sofia, Lillit, Anna T, Daniel B, Dmytro, Tamas, Jonas, Kerstin, Arnie, Cristina, Mohammad, Maria, Vivika, Lucas, Daniel J, Malin, Lovisa, Hanfeng, Erika, Akib, Rinni, Zunyuan, Qingyan, Dogan, Alexander, Caio, Henning, Adam, Majken, Erik, Chris, Sarah, Greta, Suna, Madeleine and Jakob. As you have noticed, I strongly prefer working from the office to working from home, which is why I have had the opportunity to meet so many of you. I have enjoyed your company at many lunches, coffee breaks, beers, conferences, badminton games, runs, board games and parties. A special thank you goes to Melinda, who, besides being a great friend, also helped me through the trials of the job market. Claudia and Xiao have livened up my otherwise lonely office and been extremely tolerant of my many headphone-free Zoom meetings.

During these five years I have shared many wonderful experiences with my non-Ekonomikum friends, who have done their best to remind me that a life outside of work also exists. Thank you, Lucas, Thomas, Benjamin, Hanna, Ullis, Johan, Kathryn, Jocke, John, Niels, Uliana, Jakob, Lubomyr and Anastasiia! It has been a delight to have you in my life. I have not met you as much as I would have liked to during these years, as this thesis has consumed much of my time. I promise to improve now that I am finished with it, although the work culture in academia means that I probably still will not be able to hang out with you often enough.

My grandparents, cousins and other relatives in Ukraine provided me with a second home when I came to visit for a month every summer. Meeting them and taking a break from coding and typing to walk around Lviv, feed chickens, mow grass with a scythe or go treasure hunting in mosquito-infested forests provided an indispensable reprieve. It goes without saying that I cannot wait for the war to end to be able to visit them again.

Finally, I must thank the two people who have influenced and shaped me the most. My parents, Maria and Ihor, have cared for me and guided me since I was born. They helped me through school and inspired me to follow them in obtaining a PhD. I may have strayed from their preferences for biochemistry, but I know that they are happy regardless. All the thanks in the world for your massive support!

After having spent the past ten years at Ekonomikum, it will soon be time for me to leave its storied hallways. However, I will not be going very far. Please come visit at IBF in the fall!

Uppsala, August 15th, 2022
Yaroslav Yakymovych

Contents

Abbreviations

ATE	Average treatment effect (often in the context of *subsample* average treatment effects)
CATE	Conditional average treatment effect (often in the context of GRF treatment effect estimates)
GRF	Generalised random forest
ISCO	International Standard Classification of Occupations
OECD	Organisation for Economic Co-operation and Development
SSYK	Standard för svensk yrkesklassificering (Swedish system for classifying occupations)
WSS	Wage structure statistics

Introduction

Human labour is probably the most important resource traded in the economy. Almost all of us participate in the labour market for a large part of our lives, with working-age individuals in developed countries spending almost a quarter of their waking hours on the job (OECD, 2022). The vast majority of adults rely on labour or labour-related transfers such as pensions as their main source of income (Piketty et al., 2018; SCB, 2022). Unsurprisingly, the labour market arouses strong opinions and passions among both policymakers and society at large. Unemployment, wage differences between workers, and division of firms' revenues between capital and labour are issues that have ignited a plethora of political movements in the past, and continue to be central in today's political debates. This thesis, which consists of four self-contained essays, aims to improve the understanding of several key aspects of the labour market from a worker perspective.

Technological Change and the Nature of Work

The way we work has long shifted with changes in society and technology. This process has been especially visible during the last two centuries or so, as the pace of economic growth has increased exponentially with the onset of the Industrial Revolution. The initial introduction of steam engines and textile-processing machines has been followed by railroads, electrification, internal-combustion engines, communications technology and computers. In recent decades, digitalisation and globalised trade configurations have had large effects on working patterns. While technological progress has in general been rapid and led to hitherto undreamed-of improvements in living standards, not everyone has been able to reap the gains. This is encapsulated in Schumpeter's (1942) classic description of technological change as a process of creative destruction, where new patterns of work and production continuously supersede old ones, leaving those unable to adjust behind. Given the uneven distribution of returns, it is unsurprising that some oppose the disruption wrought by new technologies. Understanding the effects of technological development on its losers is key for designing policies that can help everyone enjoy its benefits.

Automation of tasks previously performed by human workers has been a feature of the economy at least since the textile-spinning machines that ushered in the Industrial Revolution at the end of the 18th century. The workers initially exposed to automation were skilled artisans with significant training, which was rendered obsolete with the invention of machines that performed the same tasks several times as fast. This gave rise to early protests against technological change in the form of the Luddite movement, whose adherents smashed the industrial equipment that had stolen their livelihoods. Since the middle of the 20th century, however, automation has instead targeted tasks that used to be performed by low-educated labour. Recent advances in digitalisation and artificial intelligence also increasingly put white-collar clerical workers at risk. Some scholars and entrepreneurs warn that coming advances in these fields will lead to mass unemployment, as more and more occupations will be replaced by new technology (see, e.g., Brynjolfsson and McAfee, 2014). Most economists who have studied the issue take a more nuanced view, noting that recent automation often complements and increases the demand for skilled labour (Acemoglu and Restrepo, 2019). An interesting – and perhaps worrying – fact about recent technological advances is that they seem to have benefitted not only skilled occupations, but also the least-paid elementary and service jobs, leading to *polarisation* in the occupational structure. This has been to the detriment of *routine* jobs, which are easily replaced by technology and typically involve manufacturing or clerical work in the middle of the occupational wage distribution (Acemoglu and Autor, 2011). In Sweden, structural change has led to high-skilled professionals' share of the workforce increasing by a third between 1996 and 2013, while machine operators' and crafts workers' share fell by a fifth over the same period. Essay IV aims to understand how such shifts in occupational structure are connected to relative wages between occupations and to changes in inequality. The findings suggest that increased labour demand in growing occupations acts to increase wage inequality; however, the worker flows into these occupations that arise in response have a counteracting effect.

Establishment Closures and Mass Layoffs

The creative destruction wrought by technological and economic change leads to a constant churn of firms, as some rise with new opportunities, while others are unable to adapt and fail. Struggling plants often close down or lay off a large share of their workforce, resulting in many individuals losing their jobs at the same time. Such mass layoff events have been the subject of a large literature in economics since the seminal study by Jacobson et al. (1993). These studies are almost unanimous in concluding that affected workers suffer significant, long-term scarring effects in terms of employment, wages and other labour market outcomes compared to similar workers who are not displaced (see Kuhn, 2002, for an overview). Job displacement is even associated with worse outcomes in terms of family formation and health (Sullivan and von Wachter, 2011; Eliason, 2012). Many papers have tried to identify groups of workers who cope relatively well with displacement and those who are severely impacted. Essay II contributes to this literature by investigating how routine workers cope with job displacement compared to non-routine workers, a question which is pertinent given that routine workers have suffered from worsening employment opportunities and slow wage growth at least since the 1980s (Cortes, 2016). As expected from theory, routine workers are much more adversely affected by job loss, with worse labour market outcomes compared to non-routine workers for many years after displacement. In Essay III, we take a broader view and analyse how a large array of worker characteristics combine to determine the size of displacement losses. The findings point to age and education as key correlates of post-layoff outcomes, with older and low-educated workers losing more. Industry and location-related factors also play an important part. However, no single variable or small number of variables can explain all of the variation in the impact of job displacement. This confirms the findings of previous studies, which have identified many drivers of heterogeneity in this setting.

Sickness Insurance and Sickness Absence

Job displacement due to establishment shutdowns or mass layoffs is not the only kind of negative shock workers might experience. A very common negative event, which influences labour market productivity, is a deterioration in health. Everyone gets ill from time to time, being unable to carry out their usual tasks at work. Sickness insurance systems aim to compensate individuals for the income losses that arise when they are unable to perform their jobs due to their health status. Comprehensive systems, such as those found in most European countries, cover workers in all sectors of the economy and often reimburse a large fraction of labour income (Barmby, 2002). However, the design of sickness insurance must take into account the moral hazard problem inherent in all insurance settings. This is the phenomenon of insurance recipients using their information advantage over the insurer to extract more than their fair reimbursement. In the case of sickness insurance, moral hazard can involve workers taking out more days of sickness absence than their health status warrants. Policymakers employ a variety of measures to reduce such risks, including not reimbursing all lost income, excluding the first day of sickness absence from insurance coverage, and limiting the duration of benefits. Essay I studies the very common measure of using medical professionals to monitor absent workers' health status, focusing on identifying groups of workers whose behaviour changes strongly in response. The results suggest those with a history of high sickness absence uptake, men, low-income workers, and residents of socioeconomically disadvantaged neighbourhoods are more prone to increase absence when monitored less. My general approach, which involves analysing many characteristics, is able to identify the key importance of sickness absence history and neighbourhood factors, which have not been emphasised in previous literature on monitoring sensitivity.

Econometric Methodology

Methodological advances in econometrics have been key to improving our understanding of different aspects of the labour market in recent decades. The “credibility revolution”, pioneered in the 1990s by the 2021 Nobel Prize winners Joshua D. Angrist, David Card and Guido W. Imbens, has led to higher standards in research design, shifting focus to experiment-like settings where causal effects can be discerned given reasonable assumptions. While economists are sometimes lucky enough to be able to conduct or analyse randomised experiments (as I am in Essay I), we typically have to make do with observational data. However, in many observational settings, as-if-random allocation of individuals to different treatments can arise even in the absence of any conscious randomisation. Essays II and III make use of such an allocation of workers into job loss due to mass layoffs. One upside of using mass layoffs to understand the consequences of job loss is that those who are displaced when an establishment closes are less selected on their individual characteristics than those who are displaced in smaller layoffs or fired for personal reasons. Therefore, based on a broad set of worker, establishment and aggregate characteristics, it is possible to plausibly identify a control group of non-displaced workers whose characteristics closely match those of the displaced workers. Assuming that allocation into displacement is random, conditional on being in the matched sample, it is possible to identify causal effects on the individuals who lose their jobs. Of course, one might still be concerned about whether the variables used in the matching really capture everything that determines whether someone experiences a mass layoff or not. Nevertheless, such an approach provides substantial gains compared to purely descriptive studies.

Often, different individuals react differently to labour market shocks or policies. Besides understanding the main effect, we are frequently interested in how it varies for people with different characteristics. For example, a support programme for laid-off workers is unlikely to be socially beneficial if directed to those who are able to cope with the shock of displacement well on their own. Monitoring of sickness insurance recipients who would not be absent from work for longer than their health warrants anyway is unnecessary. For this reason, it is important to understand how effects vary across workers with different characteristics. Traditionally, such heterogeneity analysis is done by splitting the sample based on a characteristic of interest, and analysing whether treatment effects differ among individuals on either side of the threshold. This is the approach I employ in Essay II. However, when analysing heterogeneity along many different dimensions, traditional econometric methods can be insufficient. One constraint is of course time, as testing many different splits of the data can take very long. However, the conceptual problem is related to testing multiple hypotheses about the existence of heterogeneity across many thresholds. Splitting the sample along many dimensions makes finding spurious heterogeneity likely, as there is a high probability of uncovering patterns that hold within the particular sample considered by chance, but are not valid in the full population (a phenomenon known as *overfitting*). These problems can be solved by machine learning, an approach which has entered economics quite recently (Athey and Imbens, 2019). Using machine learning, it is possible to try out many ways of capturing heterogeneity in the data, and to make sure that the identified relationships are not spurious sampling artefacts. The *generalised random forest* (*GRF*; Athey et al., 2019), used in Essays I and III, can analyse how a large number of worker characteristics are connected to post-displacement earnings losses and sensitivity to sickness absence monitoring respectively. GRF splits the sample in turn at every possible threshold level of each included variable, and selects the division that maximises treatment effect heterogeneity across the resulting groups of workers. It also takes a number of measures to reduce the overfitting problem. The workers are divided into several sets and the treatment effects for workers in each set are predicted using only workers in the other sets. Hence, the predicted effect for an individual is not based on that individual's outcome, but only on the outcomes of others. This avoids fitting estimates to individual-level idiosyncrasies and ensures that the relationships identified are characteristic of the entire population.

The four self-contained essays in this thesis are described in more detail below.

The Essays

I. Who (Mis)uses the Sickness Insurance System? Evidence from a Randomised Experiment

I investigate how monitoring of sickness insurance recipients by physicians, a common measure in many countries, affects their sickness absence. The setting is a unique randomised controlled experiment, which was conducted in two regions of Sweden in 1988. During the experiment, those born on odd dates were required to provide medical certificates after the usual seven days of absence; however, those born on even dates could be away from work for an extended 14 days before having to provide a certificate. Earlier work by Hartman et al. (2013) has shown that the experiment increased sickness absence among the less-monitored group. This made it a failure in the eyes of policymakers, who decided to discontinue the trial and monitor all workers after the normal seven-day period. I use GRF to study whether it is possible to relax monitoring in a targeted way for groups of workers who do not increase their absence much in response. I find strong evidence of heterogeneity in worker responses to relaxed monitoring. Those who strongly increased sickness absence uptake had taken many days of sick leave in the past, were men, had low earnings and education, resided in socioeconomically disadvantaged neighbourhoods, and worked at large plants with high sick leave uptake among the workforce. The key importance of sick leave history and neighbourhood of residence have not been identified by previous studies, highlighting the strength of the GRF in analysing heterogeneity across many variables at the same time. Interestingly, some groups of workers with high sick leave uptake in general do not react much in terms of sickness absence when monitoring is reduced. For instance, women take out more sickness absence than men, but are not as sensitive to monitoring as men are. I argue that policymakers can make use of the findings to relax monitoring for some worker groups in order to conserve healthcare resources. In particular, my results suggest that it would be possible to let workers who do not have a history of extensive sickness absence to provide medical certificates after 14 days of absence instead of after seven days. If these workers' absenteeism increases as a response, they will eventually end up in the more stringently monitored group, making the system self-regulating.

II. Consequences of Job Loss for Routine Workers

This essay considers how routine workers fare compared to non-routine workers following establishment shutdowns or mass layoffs. I find that laid-off routine workers suffer significantly higher earnings, wage, employment and unemployment penalties. These differences persist in the medium run, with the earnings losses of routine workers remaining significantly larger than those of non-routine workers for eight years after layoff. Some of these additional losses are likely to arise because routine workers have a hard time finding new employment that fits their skills and human capital. They have a higher probability of switching both occupation and industry after becoming displaced, meaning that their prior knowledge is less useful in their new jobs. While recommendations for coping with technological change for routine workers often involve re-training and re-skilling to be competitive in non-routine jobs, I do not find indications that routine workers who switch into non-routine occupations do better than those who stay in routine ones. This is somewhat disheartening, but it might also suggest that not enough has been done to prepare these workers for a changing labour market. Perhaps tailoring policies specifically to the needs of routine workers might improve their outcomes.

III. Worker Attributes, Aggregate Conditions and the Impact of Adverse Labor Market Shocks (with Susan Athey, Lisa Simon, Oskar Nordström Skans and Johan Vikström)

We investigate how different worker and aggregate characteristics interact to determine displacement losses following mass layoffs. We take a big-picture view, considering 43 factors, which capture individuals' demographics, family situation, geographical mobility history, education, tenure, characteristics of the closing establishment, developments in their industry, features of the local labour market and the situation in the aggregate economy. We document substantial heterogeneity in earnings losses using GRF. The worst-affected decile of workers loses 46 percent of pre-displacement earnings and the most resilient decile loses only five percent. Our findings suggest that age and education are important determinants of the size of earnings losses. Older and low-educated workers suffer much more than younger and highly educated ones. Factors at the industry and location level also play an important part. Workers displaced from manufacturing industries and in rural locations tend to lose much more than their non-manufacturing or urban counterparts. However, a key result is that no single factor, or small number of factors, can fully explain heterogeneity in post-displacement earnings losses. The most resilient quartile among workers who are older than 50 and have at most vocational high school education does as well as the worst-affected quartile among workers who are younger than 30 and have at least a bachelor's degree. The fact that the determinants of earnings losses are many and complex can explain why earlier studies (e.g. Schmieder et al., 2020; Lachowska et al., 2020; Gathmann et al., 2020) have found heterogeneity in the impact of displacement across multiple dimensions. Our paper connects them by providing a detailed description of the characteristics of workers who experience large and small displacement losses.

IV. Understanding Occupational Wage Growth (with Adrian Adermon, Simon Ek and Georg Graetz)

This essay documents how the occupational and wage structures of the Swedish labour market have evolved from 1996 to 2013. During this period, the occupational structure changed significantly, while relative wages across occupations changed little, a pattern also seen in most other developed countries. This presents a puzzle because it implies that the workforce has re-sorted itself across occupations without being incentivised by a price mechanism. We find that the relative price of work has indeed increased in growing occupations, but that this is partly masked in the raw data by the effects of worker sorting. When the price of labour increases in an occupation, workers from other occupations who are less suited to performing it are drawn in. The wages of these new workers reflect their lower productivity, pushing average wages in the occupation down. When netting out composition effects, the resulting wage premia are positively correlated with occupational employment growth, providing a partial solution to the initial puzzle. However, in order to estimate the premia correctly, it is necessary to take into account the effects of lifecycle wage profiles in different occupations. We develop a method for identifying such profiles by using the flat spot in wage growth that workers experience after several decades of labour market experience. We end the paper by decomposing the increase in between-occupational wage inequality that has taken place in Sweden into factors related to wage premia and worker flows. We find that changes in relative premia have tended to increase inequality, while worker flows in response to them have worked in the opposite direction.

References

Acemoglu, Daron, and David Autor (2011). "Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics*, Vol. 4, Part B, 1043-1171. Elsevier.

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

Athey, Susan, and Guido W. Imbens (2019). "Machine Learning Methods That Economists Should Know About." *Annual Review of Economics*, 11, 685-725.

Athey, Susan, Julie Tibshirani, and Stefan Wager (2019). "Generalized Random Forests." *Ann. Statist.*, 47(2), 1148-1178.

Barmby, Tim A., Marco G. Ercolani, and John G. Treble (2002). "Sickness Absence: An International Comparison." *The Economic Journal*, 112(480), F315–F331.

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

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.

Eliason, Marcus (2012). "Lost jobs, broken marriages." *Journal of Population Economics* 25(4), 1365-1397.

Gathmann, Christina, Ines Helm, and Uta Schönberg. "Spillover effects of mass layoffs." *Journal of the European Economic Association*, 18(1), 427-468.

Hartman, Laura, Patrik Hesselius, and Per Johansson (2013). "Effects of Eligibility Screening in the Sickness Insurance: Evidence from a Field Experiment." *Labour Economics*, 20, 48-56.

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

Kuhn, Peter J., ed. (2002). *Losing work, moving on: International perspectives on worker displacement*. Kalamazoo, Michigan: W.E. Upjohn Institute for Employment Research.

Lachowska, Marta, Alexandre Mas, and Stephen A. Woodbury (2020). "Sources of Displaced Workers' Long-Term Earnings Losses." *American Economic Review*, 110(10), 3231-66.

OECD (2022). *Labour Force Statistics*, accessed on 2022-08-14.

Piketty, Thomas, Emmanuel Saez, and Gabriel Zucman (2018). "Distributional national accounts: methods and estimates for the United States." *The Quarterly Journal of Economics*, 133(2), 553-609.

SCB (2022), *Inkomststruktur efter deciler 1991–2020*, accessed on 2022-08-14.

Schmieder, Johannes F., Till M. von Wachter, and Jörg Heining (2022). *The Costs of Job Displacement over the Business Cycle and Its Sources: Evidence from Germany*. National Bureau of Economic Research, Working Paper No. w30162.

Schumpeter, Joseph (1942). *Capitalism, Socialism and Democracy*. Harper & Brothers.

Sullivan, Daniel, and Till von Wachter (2009). "Job Displacement and Mortality: An Analysis Using Administrative Data." *The Quarterly Journal of Economics*, 124 (3), 1265-1306.

Essay I. Who (Mis)uses the Sickness Insurance System? Evidence from a Randomised Experiment

I am grateful to Per Johansson and IFAU for data access and to Oskar Nordström Skans, Stefan Eriksson, Georg Graetz, Adrian Adermon, Johan Vikström, Lisa Simon, David Strömberg, Stefan Pitschner, Erica Lindahl, seminar participants at Uppsala University, IFAU, IFN and Aarhus University and conference participants at SUDSWEC 2021 for their helpful suggestions and comments. Financial support from Vetenskapsrådet, grant number 2018-04581, is gratefully acknowledged.

1. Introduction

Sick leave is a key right that is enjoyed by most workers in developed countries today, allowing them to stay home when their health is too poor to be able to work. It prevents incapacitated individuals from facing the choice of working in spite of their condition or losing their employment. Workers are shielded from the income effects of health shocks and are able to smooth their consumption over time in a way that would otherwise be unfeasible.

However, as with any insurance system, insuring workers against bad health carries a risk of moral hazard. Society has an interest in minimising overuse, both in order to guarantee proper use of public funds, as well as to ensure that the social insurance system is seen as fair and legitimate. This is especially imperative in light of the substantial public spending on sickness and disability insurance, amounting to two percent of GDP in OECD countries in 2017 (OECD, 2021). Sickness insurance recipients almost invariably have better knowledge of their health status than the insurer does; monitoring, in the form of doctor's visits or otherwise, is used to reduce this information discrepancy. However, monitoring is costly, as it typically involves engaging medical professionals, the opportunity cost of whose time is high. Therefore, focusing monitoring efforts on the most responsive individuals would improve the system's efficiency and reduce costs. Identifying heterogeneous behavioural effects when the intensity of monitoring is varied is thus of key policy relevance.

This paper investigates the sensitivity of different worker groups to monitoring using a large-scale randomised controlled experiment conducted in two Swedish regions in 1988. In the experiment, individuals were randomised into treatment and control groups based on whether they had odd or even dates of birth. Those with odd dates of birth were required to provide medical certificates if their sick leave spell was longer than seven days, while those born on even dates were only required to provide certificates if their spell exceeded 14 days. I use a machine learning approach, the causal forest (Athey et al., 2019), to identify groups of workers who are sensitive to this difference in monitoring intensity. Causal forests have been designed specifically for the study of heterogeneity in treatment effects. In contrast to traditional sample splitting approaches, it is the algorithm itself that chooses what characteristics to split on and at what thresholds. This minimises the researcher’s ability to select splits that fit a particular hypothesis. Furthermore, the causal forest provides individualised estimates of treatment effects for every sickness insurance recipient based on his or her specific combination of characteristics. This makes it possible to describe what characterises individuals who are responsive to reductions in monitoring. The causal forest is also able to capture complex nonlinear relationships between worker characteristics and monitoring sensitivity, as well as interactions between characteristics, in a very flexible way.

I identify substantial heterogeneity of worker responsiveness in terms of sickness absence duration. The least sensitive decile of workers is estimated to increase the duration of their sick leave spells by 0.36 days on average, while the most sensitive decile's sick leave spell duration increases by 1.71 days. The probability that a worker in the most sensitive decile ends his or her sick leave spell in its second week is estimated to increase by 19 percentage points, compared to 6 percentage points for a worker in the least sensitive decile. The most important predictors of strong worker responsiveness to monitoring are a history of high sick leave absence, low socioeconomic status in terms of education and income, male gender, large workplace, high sick leave take-up by colleagues at the workplace in earlier periods, and low socioeconomic status of the neighbourhood of residence. For predicting which workers change their sickness absence behaviour in response to monitoring, sick leave history is particularly important, as it is observed by the insurer and thus feasible to use as grounds for focusing monitoring efforts. The results regarding peer effects indicate that behaviour of colleagues is important and that interventions by firms to improve the work environment or morale can be useful.

Back-of-the-envelope calculations suggest that if monitoring intensity is reduced for workers who are estimated to be non-sensitive, rather than for a random subset of workers, losses in terms of increased sickness absence can be limited. Relaxing the monitoring regime for the 49 percent of workers with the smallest predicted treatment effects (rather than for a randomly selected 49 percent of workers as in the experiment) would result in absence rising by 41 percent less than was observed. If monitoring could only be targeted based on workers' sick leave history, absence would still increase by 29 percent less than what was the case when monitoring was randomly assigned.

Much of the literature on sickness absence has focused on identifying worker characteristics which are correlates of sickness absence uptake (Winkelmann, 1999; Barmby et al., 2002; Frick and Malo, 2008; Treble and Barmby, 2011). It is well established that sickness absence is higher among women, public sector employees, low-paid workers, high-tenured workers, and employees at large workplaces. While I find that some factors associated with high uptake, such as low earnings and working at a large establishment, are correlated with high sensitivity to monitoring, this is not the case for a number of other covariates. In particular, women and public sector workers are less responsive to monitoring than men and private sector workers. Some other correlates of sick leave uptake, such as age, marital status and workplace tenure do not have strong relationships with monitoring sensitivity.

Examples of more direct studies of sensitivity to monitoring are Ferman et al. (2021) and Boeri et al. (2021). Both have focused on public sector workers, using a policy change in a Norwegian municipality and an experiment in Italy respectively. The studies arrive at opposing conclusions, with Ferman et al.

(2021) finding no increase in sickness absence when medical certificate requirements are relaxed, and Boeri et al. (2021) finding that random visits to the homes of absent employees do have an absence-reducing effect.

Earlier work on the 1988 Swedish monitoring experiment has found that the relaxed rules caused a substantial increase in sick leave spell duration among treated workers. If applied nationally, the less stringent rules would result in costs rising by about one billion SEK (200 million 2021 EUR), representing three percent of the outlays of the entire sickness insurance system (Riksförsäkringsverket, 1989; Hartman et al., 2013). Some previous studies have investigated how the experiment's effects differed based on worker characteristics such as age, gender and income (Hesselius et al., 2009; Hartman et al., 2013; Hesselius et al., 2013; Johansson et al., 2019). In this paper, the question is approached in a more general way, by considering heterogeneity across a total of 40 individual characteristics. Furthermore, the causal forest represents a methodological advance over traditional sample splitting, as outlined above. The findings of earlier studies regarding higher sensitivity among men and individuals with low incomes are confirmed by my analysis. However, the key importance of sick leave history, as well as of neighbourhood characteristics, has not been identified previously.

There are large and persistent cross-country discrepancies in levels of sick leave take-up, the reasons for which have not been conclusively identified. Explanations put forward include differences in monitoring intensity, replacement rates, workforce health, as well as cultural factors (Barmby et al., 2002). Evidence from Sweden suggests that lowering replacement rates and excluding the first day of sickness absence from insurance coverage reduces absence rates (Johansson and Palme, 2002; Henrekson and Persson, 2004). Although differences exist, Sweden's public corporatist sickness insurance system is broadly similar to those found in the rest of Scandinavia and continental Europe. Provisions with regard to high replacement rates, lack of an unpaid initial waiting period and lack of monitoring for short absence spells were quite generous at the time of the experiment, but they are not unlike those prevailing in other European countries today (Palme and Persson, 2020). The results thus provide an insight into how monitoring affects recipient behaviour in an institutional setting characteristic of many developed countries.

The issue of monitoring sickness insurance recipients rose to renewed prominence during the Covid-19 pandemic. Many countries relaxed rules for obtaining sick leave (OECD 2020). For example, in Sweden, the maximum period a worker could spend on sick leave before having to provide a doctor's certificate was increased from 7 to 21 days during many phases of the pandemic (Försäkringskassan, 2021). This relaxation of monitoring intensity was very similar in spirit to the changes effected by the 1988 experiment.

The remainder of this paper is structured as follows. Section 2 provides background about the Swedish sickness insurance system and the context in which the monitoring experiment took place. The outcomes, as well as worker characteristics considered as possible drivers of treatment effect heterogeneity are covered in Section 3. An overview of the machine learning approach used to identify conditional treatment effects is given in Section 4. Section 5 provides evidence that the experimental randomisation was successful and Section 6 presents and discusses the results. Section 7 concludes.

2. Background

2.1 The Swedish Sickness Insurance System

Sweden has a comprehensive sickness insurance system, where practically all employees are entitled to sick leave. In 1988, at the time of the experiment, government-run social insurance funds, each responsible for a certain geographic area, covered the vast majority of sick leave expenses. Recipients were entitled to payments from these funds starting from the first day of absence. Normally, workers were required to provide the insurance funds medical certificates proving that they were sick if the duration of absence was eight days or longer. Workers were reimbursed 90 percent of their wages while they were on sick leave (SOU 1981:22).[1] However, benefits were capped for workers whose annual earnings were greater than 193,500 SEK in 1988 (equal to about 69,000 2020US$). About 2.6 percent of the workers involved in the monitoring experiment had earnings in excess of this cap. There was no time limit on benefit duration. Leave for taking care of sick children was, and has remained, separate from sickness absence in the Swedish system (SOU 2015:21). The rules for such leave were unaffected in the 1988 experiment (Riksförsäkringsverket, 1989).

Since 1988, a number of changes have been enacted to the system, mostly with the aim of reducing moral hazard and overuse. Replacement rates have been reduced to 80 percent of wages, limits on the maximum duration of sickness absence benefits have been introduced, recipients are no longer reimbursed for the first day of sick leave (the "qualification day") and the first two weeks of sick pay are now paid by employers rather than the public insurance system (SOU 2015:21). However, the Covid-19 pandemic brought about a loosening of some rules. Most notably, the qualification day rule was not applied and recipients were reimbursed for all their days of sick leave. Also, monitoring in the form of medical certificate requirements, which had been

[1] For some workers, such as municipally employed workers and white-collar workers in many collective agreement fields, there was an additional amount paid by their unions. For these individuals, the replacement rate of sickness insurance could amount to 100 percent for short to intermediate length spells. This also mitigated losses for those with earnings in excess of the reimbursement cap (SOU 1981:22). The rules for providing medical certificates for these additional reimbursements were also changed in line with the experiment (Riksförsäkringsverket, 1989), meaning that affected individuals faced no asymmetric incentives.

required from the eighth day of absence on, only took place from the 21st day of absence on (Försäkringskassan, 2021).

2.2 The “Extended Right to Self-Accorded Sickness Absence” Experiment

In 1984, the Swedish government implemented a trial known as the “Free Municipality Experiment”, which meant that a number of municipalities gained the right to try out new policies regarding, among other fields, healthcare, schools, social planning, labour market policy and environmental policy (SOU 1991:68). Within the framework of this experiment, Jämtland county in northern Sweden implemented a policy known as “Extended Right to Self-Accorded Sickness Absence” starting on January 1st 1987. This involved extending the time of sickness absence that individuals could take out without providing the government-run insurance society with a medical certificate from 7 to 14 days. This involved all workers, without the date of birth differences that were introduced by the experiment later. The motivation was that examining workers on sick leave and writing certificates for them was a waste of doctors’ time, which could be better spent treating seriously sick patients. Another motivation was that sick individuals who might find doctor’s visits a nuisance would stay at home until full recovery, thus improving their future health status and reducing future sickness absence. Also, workers would return to work when they felt well enough to do so, rather than waiting for the full number of days that the doctor had specified on the medical certificate; thus, some spells were actually expected to become shorter. A final reason for the change in policy was that travel distances to the nearest medical establishment can be very large in rural Jämtland, placing an undue burden on recipients (Riksförsäkringsverket, 1989).

The local authorities in Jämtland considered the new policy successful, but the Central Insurance Agency wanted a more rigorous evaluation, also involving a prominent urban area, as sickness absence in Sweden was higher among urban workers at the time. Thus, a randomized experiment was set up involving the 70,000 sickness insurance recipients in Jämtland county and the 240,000 recipients in Sweden’s second largest city of Gothenburg. Those born on odd dates were required to provide doctor’s certificates starting on the eighth day of their absence spell, while those born on even dates were required to provide certificates starting on the fifteenth day. Thus, the experiment represented a loosening of the rules in Gothenburg and a tightening of the rules in Jämtland. Nevertheless, throughout this paper, I refer to those who had to provide certificates on day eight as the *control group* and those who had to provide certificates on day fifteen as the *treated group*. The experimental rules covered sickness absence spells which began between July 1st and December

31st 1988. There was a substantial information campaign targeted at the insured regarding the experiment, involving leaflets distributed at workplaces and articles in the press. Subsequently, evaluators at the Central Insurance Agency assessed recipients' understanding of the experimental rules as very good, although there were a few isolated misunderstandings involving individuals in the control group thinking that looser treatment group rules were applicable to them (Riksförsäkringsverket, 1989).

Already at the preliminary stage of collecting results, there were strong indications that the length of the treated group's sickness absence spells had increased substantially and the experiment was discontinued, with everyone in both Jämtland and Gothenburg having to provide doctor's certificates from the eighth day for spells which began on January 1st 1989 or later. Evaluators at the Central Insurance Agency (Riksförsäkringsverket, 1989) later estimated that the less stringent rules would lead to an increase in costs of about 1 billion SEK (some 3 percent of the total costs for the entire sickness insurance system) if applied nationally. The findings of Hartman et al. (2013) confirm this, showing substantially longer absence duration for the treated group, with sharp changes in spell survival and hazard rates indicative of moral hazard.

Some groups of workers were excluded from the experiment for administrative reasons. The largest of these, some 11 percent of the workforce, were individuals whose employment contracts were regulated by the central government, including teachers, postal workers, government agency employees, railway employees, police, military servicemen, sickness insurance fund employees, customs and border guards, government-owned forestry company workers, Church of Sweden clergy, university employees, and others. The reason for the exclusion was that sick pay to these workers was provided directly by their employers, who then in turn were reimbursed by the social insurance funds. There was also a very small group of individuals who were required to provide medical certificates already on the first day of sickness absence, mostly due to prior misuse, to whom the experiment did not apply (Riksförsäkringsverket, 1989).

3. Outcomes and Characteristics

Thanks to unusually rich microdata collected by Statistics Sweden, I am able to include a broad set of worker characteristics in the analysis. These contain information on sickness absence (starting in 1987), demographic characteristics, place of residence, employment relationships and earnings (starting in 1985) and family situation (imputed from 1990 data).

3.1 Outcome Definitions

Earlier work on the experiment by Hartman et al. (2013) has found sizeable effects on the duration of treated workers' sickness absence spells, but no evidence of an effect on sickness spell incidence. Because of this, my analysis focuses on the intensive rather than the extensive margin.

3.1.1 Duration of Sickness Absence Spell

The main outcome studied is the duration of spells of sickness absence in days. This is a natural margin to consider, as sickness insurance costs scale with absence duration. There were a total of 256,465 sickness spells started by individuals in the studied sample between July 1st and December 31st, 1988. I drop spells whose duration makes it unlikely that they were affected by differences in monitoring between days 7 and 14. The survival and hazard graphs strongly suggest no differences between treated and controls for spells shorter than four and longer than 21 days. These two categories comprise 112,467 and 19,796 spells respectively. Including long spells is problematic, as they might have outsize effects on estimates due to being numerical outliers. The very large number of short spells would also serve to obfuscate patterns of behaviour during the period when monitoring intensity varied. For this reason, only spells between four and 21 days in duration are used in the main analysis.

3.1.2 Probability of Sickness Absence Spell Lasting 8-14 Days

Another way of measuring individual responsiveness to the experiment is by studying the probability of a sickness absence spell ending during its second week. For this outcome, spells shorter than four and longer than 21 days are retained, as outlier effects are absent due to its binary nature. The results thus serve as robustness tests for both the outcome definition as well as for the sample restrictions imposed in the main analysis. As explained in Section VI, the two set-ups produce qualitatively similar findings.

3.2 Worker Characteristics

The selection of worker characteristics included in the analysis is based on factors which have been identified as important for sick leave uptake by previous literature. Most of these factors have not been linked to monitoring sensitivity. Nevertheless, a simple hypothesis would be that groups with high uptake also respond more strongly to being monitored.

3.2.1 Health-Related

An individual's health status is, in the absence of moral hazard, the only determinant of sick leave duration. Unfortunately, the necessary data for fully characterising individuals' health are lacking. Information about inpatient care spells and diagnoses received during such spells are only available starting on

January 1st, 1987. Data on outpatient care contacts and diagnoses, which constitute the vast majority of medical treatment in Sweden, are unavailable for the period studied. For this reason, the two measures of individual health I use are indirect. The first of these, *the total number of days of sickness absence in earlier periods,* is the total number of days of sickness absence the individual ran up in spells which began between January 1st, 1987 and June 30th, 1988.[2] This measure contains information not only on the individual's health, but also on any overuse of sickness insurance that the individual might have been prone to. A measure much more directly related to serious health issues is *the total number of days spent in inpatient care* over the same time frame.'

Another variable connected both to health and to sickness absence behaviour is *the number of short sickness absence spells in earlier periods*. Short spells are defined as those 1-21 days in length. This measure puts less weight on long spells, instead focusing on whether the individual has taken many short absences in the past, which might be indicative of the presence of moral hazard. The number of spells is also measured between January 1st, 1987 and June 30th, 1988.

3.2.2. Demographic

There are well-established differences between demographic groups in terms of their sick leave uptake. For example, a *female* dummy is included because women have a persistently higher take-up than men. Also, health deteriorates with *age*, which has been shown to affect sick leave (Barmby et al., 2002). Finally, in the Swedish setting, immigrants tend to be absent due to sickness slightly more than natives; this factor is captured using an *immigrant* variable which takes the value 0 for individuals born in Sweden, 1 for those born in other Nordic countries, 2 for those born in the rest of Europe and 3 for those born in the rest of the world. The causal forest results are interpretable even though the measure is ordinal, which would not be the case for traditional econometric approaches.

3.2.3 Family-Related

Family factors have been found to play a role in workers' sickness insurance take-up. The presence of partners may affect behaviour through their provision of additional sources of income and married individuals tend to have higher sick leave uptake than unmarried ones (Barmby et al., 2002; Angelov et al., 2011). To analyse the importance of such effects, dummies are included for being *married* and being *divorced*, with single individuals providing the reference category. While leave for taking care of sick children is separate from sick leave in Sweden, parents might nevertheless register such spells as

[2] The length of the pre-period considered is dictated by data on sickness absence becoming available from January 1st, 1987.

own sickness absence. For this reason, I include the *number of children of pre-school age* and the *number of children of school age* as characteristics. As family characteristics are only available starting in 1990, I impute their 1988 values. Individuals' marital status is imputed to be the same in 1988 as in 1990; two years are subtracted from the ages of children in 1990 to get values for 1988.

3.2.4 Education

Education is a strong correlate of factors identified as important for sick leave uptake, such as earnings and occupation. To flexibly capture education, I have included both the *level of education*, as an ordinal variable running from 1 (mandatory 9-year education or less) to 7 (PhD-level education), as well as *dummies for broad education fields*. The fields are *general education* (found at the low levels of educational attainment), *teacher training*, *administration/law/social science*, *science/engineering*, *health* and *services*.

3.2.5 Neighbourhood Characteristics

Individuals' sickness absence behaviour might be affected by the attitudes of their neighbours (Lindbeck et al., 2016). For this reason, I include several leave-one-out characteristics of the neighbourhood where the sickness insurance recipient lives. The neighbourhoods (called SAMS) are small, corresponding to several urban blocks or small portions of the countryside. The median number of inhabitants aged 16-64 in each neighbourhood is 398, with the mean being 586. Neighbourhood characteristics included are *average annual earnings*, *share of inhabitants with a post-secondary education* and the *immigrant share*. These three measures are constructed based on the population aged 30-64, not taking into account those past working age, or those who are likely to not have completed their education.

3.2.6 Career-Related

High-earning individuals are on sick leave less than their lower-earning peers. This could be due to better health, stronger intrinsic motivation, as well as lower income replacement rates, as is the case for individuals with earnings in excess of the replacement cap (Barmby et al., 2002).[3] These effects are captured by an *annual labour income* variable. A related concept is the worker's *income rank at his or her workplace*. The rank is measured in relative terms, with 0 representing the worker who earns least and 1 the worker who earns most regardless of workplace size. This measure also captures key worker ef-

[3] The annual earnings level of 2.6 percent of the individuals involved in the experiment was such that their replacement rates were lower than the otherwise stipulated 90 percent. At least some of these workers were however likely to be (partly) reimbursed for this loss by additional union-negotiated sickness insurance. For evidence on sickness absence being affected by replacement rates in the Swedish setting, see Johansson and Palme (2005).

fects, which imply that workers who are more important for workplace functioning are less likely to make use of the sickness insurance. This could be because they continue working even when their health status is bad, or because individuals with better average health select into such roles (Hensvik and Rosenqvist, 2019). Workplace *tenure* has been identified as a correlate of sickness absence in the literature, with tendencies for high-tenured workers to take out more sick leave than lower-tenured ones (Barmby et al., 2002). Tenure is also correlated with job security, which has been suggested to increase sickness absence (Bratberg and Monstad, 2015). The tenure measure goes from 0 to 3 years and is censored at the top because matched employer-employee data only become available from 1985 onwards.

3.2.7 Workplace-Related

Different sectors of the economy have traditionally experienced different sick leave rates (Barmby et al., 2002). This could be due to intrinsic differences in workforce characteristics, such as gender and age composition, differences in work environment quality, which cause ill health among the employees, or because some sectors are more permissive of overuse of sickness absence. The public sector has seen higher sickness absence rates than the private sector in many countries (Frick and Malo, 2008); for this reason, a *local government sector* dummy is included in the analysis, with the private sector being the baseline. This dummy takes on the value one for individuals employed at the municipality or county level. In 1988, the Swedish local government sector included healthcare, elderly care, municipal services and administrative staff. As central government employees were excluded from the randomisation, they are dropped from this study.

Differences between sectors are further captured by nine broad industry dummies: *primary*, *manufacturing*, *construction*, *utilities*, *wholesale and retail*, *business services*, *health*, *education* and *public administration*.

The *number of workers at an establishment* has been suggested to affect sick leave uptake. This could be both because large workplaces are worse for employees' health and because the importance of a single individual decreases with workplace size, meaning that costs of unnecessary absence spells are lower (Winkelmann, 1999, Lindgren 2012).

Finally, I include a dummy for whether an individual *commutes to another municipality* for work. This is to capture higher costs of getting to work, which might induce individuals to stay at home (van Ommeren and Gutiérrez-i-Puigarnau, 2011).

3.2.8 Peer Effects

I consider two kinds of peer effects. The first is the behaviour of colleagues at the individual's place of work. This is measured as the *leave-one-out average number of sickness absence days per worker at the workplace* between January 1987 and June 1988. The second peer effect relates to the sick leave behaviour of neighbours, and is computed as the *leave-one-out average number of sickness absence days among employed individuals in the neighbourhood* between January 1987 and June 1988. This measure is based on those aged 30-64 to be in line with the other neighbourhood measures.

3.2.9 Other Variables

The *population* density of *the municipality where the individual resides* is intended to capture any differences between areas with different levels of urbanisation. Relative sick leave uptake between urban and rural areas in Sweden has varied over time and even reversed. Another reason for including population density is that it captures some of the differences in travel distance to the nearest medical facility, which can be large in rural areas. A separate *dummy for Gothenburg* is included because the intensity of monitoring was pushed in different directions in Gothenburg and in Jämtland. It is not clear that the effects of a reduction in monitoring intensity should be the same as the effects of a corresponding increase in monitoring intensity. However, the two measures are highly correlated, as Gothenburg had a much higher population density (963 people/km^2) than any municipality in Jämtland (at most 26 people/km^2). Combined with the fact that 79 percent of the affected workers lived in Gothenburg, this means that the two characteristics are difficult to disentangle and their effects are considered in combination.

4. Empirical Approach

If experimental randomisation holds, the average effect of reduced monitoring for all workers can be estimated by a simple regression of absence duration on treatment. The goal of this paper is to go beyond this by estimating the effects of monitoring on the sickness absence behaviour of different groups of workers, as characterised by their attributes $\boldsymbol{x}$. Two approaches are traditionally employed for such heterogeneity analysis. The first relies on splitting the sample of workers at a threshold value $\bar{\bar{x}}$ of a characteristic x that is of interest and comparing differences between treated and controls on either side of the threshold. Alternatively, matching compares the outcome of a treated worker to outcomes of control workers based on their characteristics and a kernel which determines which control(s) the treated worker is compared to.

The causal forest instead relies on the output of a large number of causal tree algorithms, developed by Athey and Imbens (2016). Each of the causal trees splits the sample of workers into two groups based in turn on all the possible threshold levels $\bar{\bar{x}}$ of each of the analysed characteristics x. For every one of these potentially very many splits, the algorithm computes treatment effects within the two resulting worker groups. The split that yields the largest difference in estimated treatment effects is chosen. The two resulting groups of workers are recursively split again according to the same criterion. The workers are thus sorted, based on $\boldsymbol{x}$, into "leaves" with similar estimated treatment effects.

While a single causal tree finds the best fit for treatment effect heterogeneity among the sample considered, the estimates of single trees can be non-robust, and their standard errors are difficult to estimate. For this reason, a causal forest (Athey et al., 2019), which is a large collection of causal trees, is estimated instead. Each of the trees is estimated on a random subsample of the workers, meaning that they will not be identical to each other. The causal forest thus reveals relationships that hold consistently across random subsamples of workers. The causal forest's output consists of treatment effect estimates for each worker based on that worker's combination of characteristics $\boldsymbol{x}$. The causal forest can be seen as providing a highly flexible matching kernel that is determined by how often a worker appears in the same leaf of a tree as workers with characteristics $\boldsymbol{x}$.

One key advantage of causal forests relative to traditional sample splitting approaches is the data-driven way in which characteristics and thresholds are selected. A horserace is run between many different candidate splits, and the one that can explain the largest share of treatment effect heterogeneity is chosen. Unlike traditional matching, the causal forest determines both what size the relevant kernel should be, and consequently how many nearest neighbours should be included, in a data-driven way. This leaves little room for the potentially arbitrary choices that researchers might otherwise make when choosing which groups to compare.

If traditional methods are used to analyse many different splits of the data, they are likely to find spurious heterogeneity in treatment effects due to multiple hypothesis testing issues. The causal forest uses several measures to minimise this problem. When estimating treatment effects for a worker, only those of the forest's trees into which the worker was not sampled are used. This avoids fitting the estimates to the worker's own outcome. Also, in each tree, half of the selected workers are used for determining how to split the data and the other half are used when estimating effects later. This further reduces the risk of making predictions based on idiosyncratic patterns in the data.

Compared to other approaches, the causal forest can be more flexible in capturing nonlinear effects of the characteristics $\boldsymbol{x}$, as well as (potentially complex) interactions between these. This is because of its nonparametric nature, which can be thought of as fitting a step function instead of a standard regression. The causal forest as implemented in this paper has been shown to perform competitively compared to other machine learning methods by Knaus et al. (2021).

4.1 Methodological Details

If one considers the whole set of workers $\mathbb{S}$, the average treatment effect τ can be estimated by applying the following simple regression model:

$$y_i = \alpha + \tau W_w + \varepsilon_i$$

Where y_i indexes the duration of sickness absence spell i, $W_w \in \{1,0\}$ is the treatment status of the worker w, and unconfoundedness, $E(\varepsilon_i|W_w) = 0$, is assumed to hold due to experimental randomisation.

Within the potential outcomes framework, each absence spell i can be thought of as having two durations depending on the worker's treatment status, $y_i|W_w = 1$ and $y_i|W_w = 0$, written as y_{i1} and y_{i0} for simplicity. Ideally, one would want to estimate treatment effects $\tau_w = E(y_{i1}) - E(y_{i0})$ for each worker in the sample. However, this is unfeasible, as only one set of y_{i1} and y_{i0} are observed for any individual. It is nevertheless possible to estimate heterogeneous treatment effects τ_x for worker subpopulations based on their characteristics $\boldsymbol{x}_w$:

$$\tau_x = E(y_i|\boldsymbol{x}_w, W_w = 1) - E(y_i|\boldsymbol{x}_w, W_w = 0)$$

Two assumptions are required for estimation:

(A1) Unconfoundedness, $W_w|\boldsymbol{x}_w \perp \{y_{i1}, y_{i0}\}$

(A2) Overlap, $1 > \xi \geq \Pr(W_w = w|\boldsymbol{x}_w) \geq \xi > 0,\ \forall\, \boldsymbol{x}_w, W_w \in \{1,0\}$

The unconfoundedness assumption (A1) requires that treatment assignment of workers is unrelated to potential outcomes, conditional on covariates $\boldsymbol{x}_w$. The overlap assumption (A2) requires that there is no combination of characteristics $\boldsymbol{x}_w$ among workers in the treated or control groups for which there is no corresponding subpopulation in the other group.[4] If this would be the case, the causal forest would effectively be extrapolating treatment effects for this subpopulation.

Before applying forest estimation, I split the population of workers who had registered sickness spells in the second half of 1988 into a training set containing the absence spells of 80 percent of workers and a test set containing the absence spells of the remaining 20 percent. All estimation is done on the training set, with the held-out test set used for validating the predictions.

4.1.1 Tree Algorithms

Forest estimation relies on constructing a large number of recursive tree algorithms. Regression trees (Breiman et al, 1984) divide individuals into groups with similar values of an outcome y, while causal trees (Athey and Imbens, 2016) divide individuals into groups with similar treatment effects τ. In both

[4] Indeed, the probability of treatment conditional on $\boldsymbol{x}_w$ must be bounded away from zero and one by a positive ξ.

cases, the divisions are based on a vector of individual characteristics x. A tree is grown as follows:

1. The set of workers is randomly divided into two halves, which constitute the splitting and estimation subsamples. The trees are grown using only the workers in the splitting subsample; the estimation subsample is used to populate the leaves of the tree after the splits have been made, and for calculating estimates. This is required for a property known as honesty, which ensures consistency and asymptotic normality of the forest estimates (Athey et al., 2019).
2. Consider the full set of sickness absence spells in the splitting sample. This forms the parent node $\mathbb{S}_P$.
    a. $\mathbb{S}_P$ is split into child nodes $\mathbb{S}_L$ and $\mathbb{S}_R$ in turn at every possible threshold value of each included worker characteristic x. The number of possible threshold values can be large for variables such as annual earnings, or just one for a binary variable such as gender.
    b. The criterion of interest is evaluated for each possible partition into $\mathbb{S}_L$ and $\mathbb{S}_R$. In the case of regression trees, it is heterogeneity with regard to an outcome y, while in the case of causal trees, it is heterogeneity with regard to estimated treatment effects τ. The partition which maximises the criterion is selected.
    c. Steps a-b are repeated, with $\mathbb{S}_L$ and $\mathbb{S}_R$ each being considered as parent nodes. Splitting continues until the gains of further splits are lower than a tolerance threshold. The final cells into which individuals are divided are called leaves.

4.1.2 Robinson's Transformation

Before predicting treatment effects using the causal forest, the treatment status W_w and duration of absence y_w are made orthogonal to the vector of observed worker characteristics $\boldsymbol{x}_w$. This is known as Robinson's transformation and makes the causal forest an efficient R-learner (Nie and Wager, 2019). Two separate regression forests (Breiman, 2001) are constructed to estimate the conditional propensity score $\hat{e}_x = W_w|\boldsymbol{x}_w$ and marginal response function $\widehat{m}_x = y_i|\boldsymbol{x}_w$.[5] These are based on regression trees, which divide workers into leaves with similar values of treatment status W and spell duration y based on covariates $\boldsymbol{x}$.

[5] Strictly speaking, if experimental randomization holds, estimating conditional propensity scores $\hat{e}_x = W_w|x_w$ is unnecessary, as then $\hat{e}_x = e = 0.49\ \forall\ \boldsymbol{x}$. However, to be more conservative, and to avoid some minor balancing issues as discussed in Section 5, I estimate $\hat{e}_x$ using a prediction forest. This estimation provides little gain compared to the naïve model of perfect randomisation.

Each tree is estimated on subsamples of the training data, drawn randomly without replacement; half of training set workers are drawn for each tree. These are then divided in half into splitting and estimation subsamples to ensure honesty, as explained above.

The regression trees maximise the following criterion when dividing the splitting sample:

$$\frac{N_L N_R}{N_P}(\bar{z}_L - \bar{z}_R)^2, z \in \{W, y\} \tag{1}$$

where N_P is the number of sickness absence spells in the parent node, and N_L and N_R are the number of spells in the two child nodes. The criterion aims to split workers into groups with as different $\bar{W}$ and $\bar{y}$ as possible.[6] Note that more even splits are implicitly favoured over less even ones, as indicated by the multiplication term $\frac{N_L N_R}{N_P}$. Each leaf produced by the tree is populated by estimation sample workers; these workers' values of W and y are used for making predictions in the next step. If the number of trees is large, each worker will be in the estimation sample in about a quarter of cases.

When assessing $\hat{e}_x$ and $\hat{m}_x$ at a combination of covariates $\boldsymbol{x}$, the forest "pushes" $\boldsymbol{x}$ down each tree to determine the leaf where it ends up. Each worker w is given a weight depending on how often he or she is present in the leaves where $\boldsymbol{x}$ ends up in.[7] This is analogous to nearest-neighbour matching, but with the kernels defined adaptively in a data-driven way. In a forest of B trees, with the number of workers in the appropriate leaf of each tree given by N_{l_b}, a worker's weight α_w is given by:

$$\alpha_w = \frac{1}{B}\sum_{b=1}^{B}\frac{\mathbb{1}\{w \in l_b\}}{N_{l_b}} \tag{2}$$

The estimated values of $\hat{e}_x$ and $\hat{m}_x$ for a given combination of covariates $\boldsymbol{x}$ are then calculated as:

$$\hat{e}_x = \frac{1}{N}\sum_{i=1}^{N}\alpha_w W_w \quad \text{and} \quad \hat{m}_x = \frac{1}{N}\sum_{i=1}^{N}\alpha_w y_i \tag{3}$$

In this way, predictions can be obtained for each worker's values of $\boldsymbol{x}$. Note that to receive valid predictions, only the trees where the worker was not sampled into either the splitting or the estimation subsamples can be used (out-of-bag estimation). This means that when predicting for a training set worker, the output of only about half of the trees is used in practice. The estimates are used to orthogonalise treatment status and duration of absence with regard to

[6] Technically, the forest operates on absence spells rather than workers. However, as all covariates $\boldsymbol{x}$ are constant across spells for the same worker, all of a worker's spells must end up on the same side of a partition.

[7] The worker must have been randomised into the estimation sample of that tree; otherwise, the worker's weight is zero for that tree, even if his or her covariates are such that he or she would end up in the same leaf as $\boldsymbol{x}$.

the vector of covariates to obtain $\widetilde{W}_w = W_w - \hat{e}_x$ and $\tilde{y}_i = y_i - \widehat{m}_x$. These are then used to grow the causal forest in the next stage.

4.1.3 Causal Forest

The causal forest's objective differs from that of the regression forests, as it aims to estimate conditional average treatment effects τ_x rather than to predict the value of an outcome variable. The trees thus divide the set of sickness absence spells into leaves where the workers have similar responsiveness to monitoring. Otherwise, the procedure is similar in the respect that half of the workers are sampled without replacement for each tree, and the sampled group is divided into splitting and estimation subgroups. The splitting subgroup is once again used for determining which splits the tree makes, and the estimation subgroup is used for populating the leaves.

The splits are made based on the covariate and threshold value that maximise an approximation of the following criterion:

$$\frac{N_L N_R}{N_P^2}(\hat{\tau}_L - \hat{\tau}_R)^2, \qquad \hat{\tau}_S = \frac{\sum_{i \in S} \tilde{y}_i \, \widetilde{W}_w}{\sum_{i \in S} \widetilde{W}_w^2}, \qquad S \in \{L, R\} \tag{4}$$

Note that this criterion aims to maximise the difference in treatment effects across the child nodes L and R. There is also an implicit penalty for uneven splits, as in the regression forest criterion.[8] Only a random subset of the worker characteristics $\boldsymbol{x}$ is considered when making each split of the workers into $\mathbb{S}_L$ and $\mathbb{S}_R$; this speeds up estimation, but also introduces additional randomness into tree construction.

The trees continue splitting until reaching a tolerance threshold. The leaves that this provides are populated by estimation sample workers, whose values of $\tilde{y}_i$ and $\widetilde{W}_w$ are used for estimating treatment effects. Optimally, causal forests are estimated using as many trees as there are observations in the sample. However, due to computational limitations, all estimates shown are based on a forest of 5000 trees.

When evaluating τ_x at a combination of covariates $\boldsymbol{x}$, the causal forest operates analogously to the regression forest, and "pushes" this combination down its trees to find which leaf it ends up in. Workers are provided weights α_w, the value of which depends on how often the worker is found in the leaf which contains $\boldsymbol{x}$. The causal forest's estimate of the treatment effect at a particular $\boldsymbol{x}$ is given by:

$$\hat{\tau}_x = \frac{\sum_{i=1}^{N} \alpha_w \tilde{y}_i \widetilde{W}_w}{\sum_{i=1}^{N} \alpha_w \widetilde{W}_w^2} \tag{5}$$

[8] Additionally, there is an explicit limit on how uneven splits are allowed to be. At least five percent of workers in N_P must go into either of N_L and N_R, with more uneven splits not being considered.

Estimates for the training sample workers are once again out-of-bag, making use only of those trees where the worker was not included in the splitting or estimation samples. For out-of-sample prediction, as in the case of the test sample workers, the full ensemble of trees is used.

The data contain natural clusters in the form of sickness spells taken out by the same worker. To account for this structure, worker-level clusters rather than single spells are drawn when selecting the random subsample used for each tree and when splitting into training and estimation subsamples according to the honesty procedure. Furthermore, worker clusters are reweighted when estimating $\hat{\tau}_x$ so that workers who have had different numbers of sickness spells during the experiment get equal weight. Standard errors of the predicted $\hat{\tau}_x$ are cluster-robust.

5. Randomisation and Balancing

5.1 Experimental Population and Validity of Randomisation

While date of birth considered over the entire year is correlated with many important characteristics and outcomes (see e.g. Bedard and Dhuey, 2006), having an odd or even date of birth should be random, as parents are unlikely to be able to determine exact birth timing.[9] As Swedish social insurance numbers, used for reporting sick leave, include the birth date, manipulations in response to the experiment would have been prohibitively costly. The lack of differences in the behaviour of sickness insurance recipients born on different dates in the second half of 1987 (the part of the year corresponding to the 1988 experiment) is confirmed by the graphs in Figure 1. The left panel plots the survival curve, identifying the share of sickness spells still ongoing after a certain number of days had elapsed since they began. Most sickness spells are short, with some 80 percent being over within a week. Importantly, there are no visually discernible differences between workers born on odd and even dates. A fairly sharp drop in the survival rate is evident after seven days of absence, when workers in Gothenburg (79 percent of the sample) were required to provide medical certificates. There is a smaller drop at 14 days of absence, when workers in Jämtland were required to provide certificates. These drops are confirmed by the hazard graph in the right panel of Figure 1, which shows the probability of a spell which has been going on for a given number of days ending on the next day. A sharp spike in hazard is evident after seven days of absence, and a smaller one after 14 days of absence. Furthermore, there are smaller spikes at each multiple of seven; this is because

[9] There are no laws or other policies that have a differential effect based on date of birth in Sweden, so incentives to manipulate the date outside of the setting of the 1988 experiment are lacking.

Swedish doctors tend to prescribe sick leave in full weeks (Riksförsäkringsverket, 1989). The hazard graph also shows very similar behaviour patterns for individuals with odd and even dates of birth.

Figure 1. *Survival and hazard rates for sickness absence spells taken by workers born on odd and even dates in Gothenburg and Jämtland in the second half of 1987*

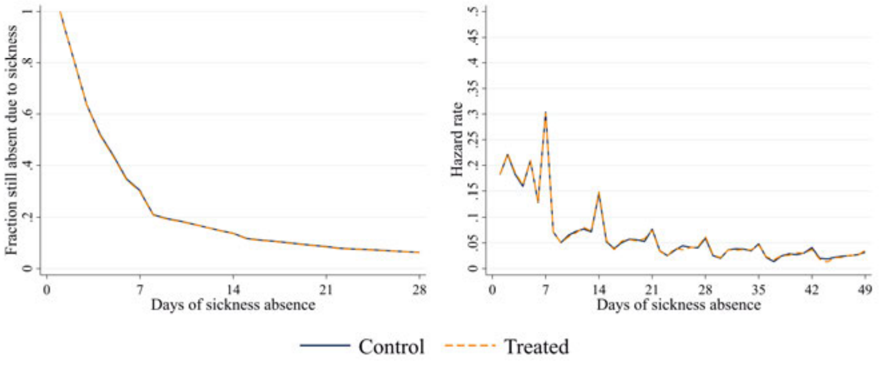

Note: Spells which began between July 1st and December 31st 1987 (pre-period). The hazard rate represents the probability that a worker who has been absent for a given number of days returns to work on the next day.

A more formal balancing table is shown in Table A1 in the Appendix. The treated and control groups are very similar with regard to all the characteristics considered. Importantly, the maximum and minimum values of each variable within the two groups align closely, suggesting that the overlap assumption is satisfied. However, a number of differences between variable means are statistically significant, even if they are numerically small. This is driven by the fact that non-European immigrants, who constitute 2.8 percent of the control group and 2.5 percent of the treated group, are more likely to be registered as born on odd dates. This is due to individuals who are not certain of their exact birth date often being registered as born on certain days of the year (commonly January 1st). This difference carries over to other characteristics correlated with being a non-European immigrant, such as neighbourhood statistics. None of the differences are likely to be economically significant, however. When non-European immigrants are removed from the sample, only the probability of holding education in administration, law and social science and the probability of commuting to a different municipality are significantly different at the five percent level; the probability of working in the education industry is significantly different at the ten percent level. This is in line with what is expected when testing for differences in means of 42 variables following a successful randomisation.

To be eligible for the experiment, workers had to be registered at the Gothenburg or Jämtland local social insurance funds and not have their wages and working conditions set by the central government. There is no exact information on dates when workers entered and exited central government employment, but I am nevertheless able to identify this group with a high degree of certainty. Individuals who were registered as central government employees in September 1988 are dropped, as are others who worked at establishments where over 90 percent of employed workers were registered as central government employees in September 1988.[10] I also exclude workers under the age of 18 and those with very low annual labour earnings (below 22,962 SEK). This leaves 77,672 workers who took out 123,429 sickness absence spells.

A final test of the randomisation is provided by treatment propensity score estimates $\hat{e}_x$ estimated on this subset of workers. The $\hat{e}_x$ are estimated by the prediction forest that is run for orthogonalising treatment status prior to estimating the causal forest. The prediction forest has many of the properties of the causal forest, including being very flexible with regard to functional form and interactions between variables when estimating $\hat{e}_x$. The same set of variables is used for predicting propensity scores as the causal forest uses for estimating heterogeneous treatment effects. Figure 2 contains histograms of these for the treated and control workers. The average propensity score is 0.49, reflecting the fact that there are slightly fewer even than odd dates. Practically all sickness insurance recipients' scores lie within 0.1 of the mean, and the distributions are very similar for both treated and controls. The prediction forest is thus unable to predict selection into treatment, in spite of its high degree of flexibility.

[10] This second restriction is to exclude workers who are likely to have been central government employees for part of the experimental period, but not in September 1988.

Figure 2. *Histogram of prediction forest treatment propensity score estimates for treated and controls*

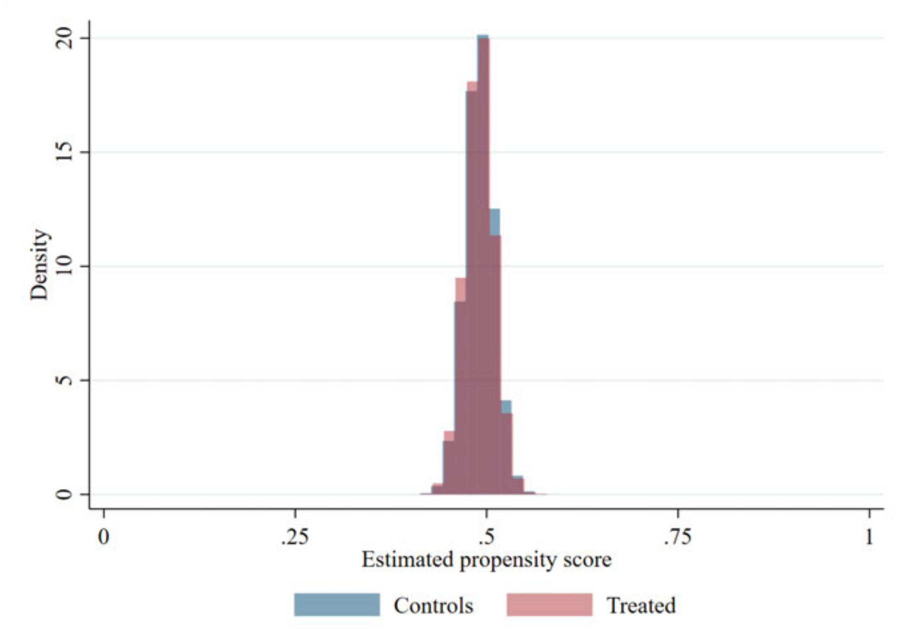

Note: Prediction forest estimates of probability of entering treatment based on the 42 included worker characteristics, plotted separately for treated and controls. Smoothed density functions overlaid.

5.2 Main Effect of the Experiment

The experiment has been evaluated by both Riksförsäkringsverket (1989) and academic literature (Hartman et al., 2013) as having a sizeable effect on the duration of sickness absence spells. Hartman et al. (2013) find that the average duration of sickness absence spells among treated workers increased by 0.6 days, but no evidence that the incidence of absence spells per worker responded to the experiment. In Figure 3, survival and hazard rates for sickness spells which began in the second half of 1988 are shown. There are striking differences in the behaviour of the treated and controls, which were absent in the pre-period (Figure 1). The survival curve for the treated is consistently above the one for the controls between days 6 and 14. The fact that the difference is present during the period when monitoring intensity differs points to the discrepancy being a result of the variation in rules. The experiment's impact is confirmed by the hazard graph, which shows large spikes in the probability of exiting sick leave at 7 days for the treated and 14 days for the controls.

Figure 3. *Survival and hazard rates for sickness absence spells taken by treated and controls workers in Gothenburg and Jämtland during the experiment*

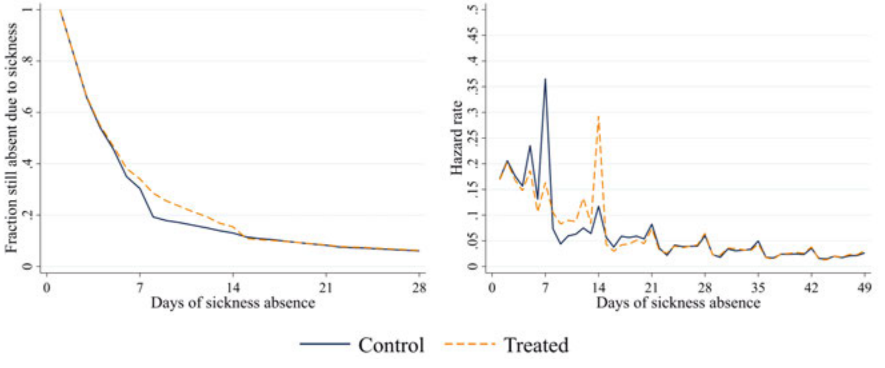

Note: Spells which began between July 1st and December 31st 1988. The hazard rate represents the probability that a worker who has been absent for a given number of days returns to work on the next day.

6. Results

6.1 Size of Heterogeneity

The distribution of estimated random forest effects is shown in Figure 4. Virtually all spells are predicted to become longer if monitoring is reduced, but there is substantial variation around the average effect of 0.9 days.[11] For the least sensitive decile, treatment effects are estimated to be at most 0.36 days, while for the most affected decile they are estimated to be 1.71 days or more. A corresponding histogram of treatment effects where the outcome is the probability of returning to work on days 8-14 is shown in Figure A1 in the Appendix. On average, treated workers are estimated to be 11 percentage points more likely to return to work in the second week of absence. This is a very sizeable effect, as the baseline probability for control workers is 8 percent. There are also large heterogeneities across groups of workers; the effects on the least sensitive decile are estimated to be less than 6 percentage points, compared to over 19 percentage points for the most sensitive decile.

[11] This figure is larger than the 0.5-0.7 days found by Hartman et al. (2013) because I drop spells shorter than four and longer than 21 days, whereas they include spells of all durations (censoring spells longer than 28 days).

Figure 4. *Distribution of predicted treatment effects on sickness absence spell duration.*

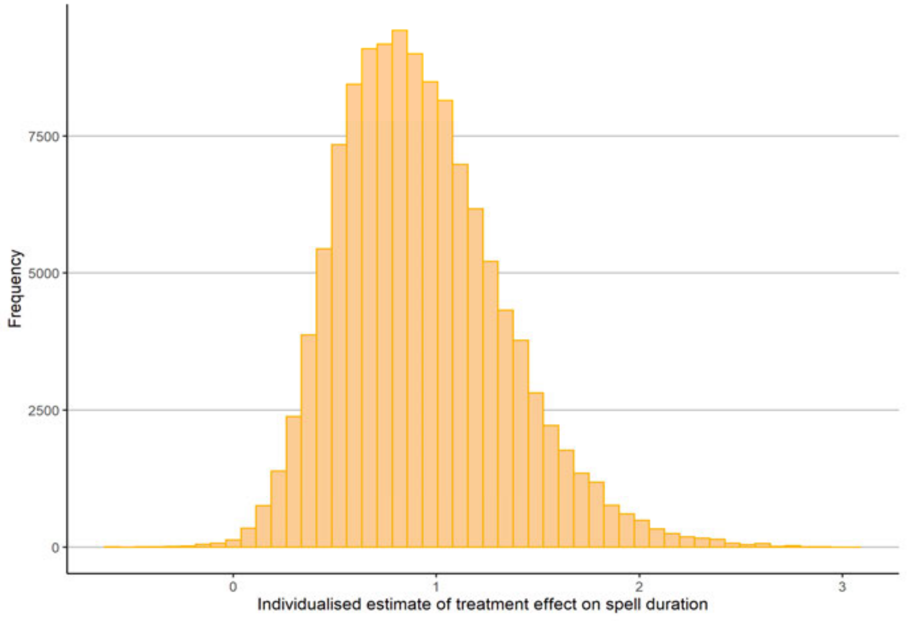

Note: Frequency represents the number of sickness spells with estimated causal forest treatment effects that fall within each bin.

The forest's performance can be assessed using an omnibus best linear predictor test (Chernozhukov et al., 2020). The best linear predictor test assesses whether both the average treatment effect and variations around this average effect are predicted correctly. The test uses the output from the prediction forests run for orthogonalisation before applying the causal forest, that is the spell durations $\tilde{y}_i = y_i - \hat{E}(y_i|\boldsymbol{x}_w)$ and treatment assignment $\widetilde{W}_w = W_w - \hat{E}(W_w|\boldsymbol{x}_w)$. The $\tilde{y}_i$ are regressed on a function of $\widetilde{W}_w$ and the causal forest's predicted treatment effects $\hat{\tau}_x$:

$$\tilde{y}_i = \alpha\left(\bar{\tau}\widetilde{W}_w\right) + \beta\left((\hat{\tau}_x - \bar{\tau})\widetilde{W}_w\right) + \varepsilon_i, \qquad \bar{\tau} = \frac{\sum_i \hat{\tau}_x}{N} \tag{6}$$

Out-of-bag estimates of $\tilde{y}_i$, $\widetilde{W}_i$ and $\hat{\tau}_x$ are used.[12] The parameter α estimates how well the forest's average predicted treatment effect fits the data. If the causal forest's out-of-bag prediction of the average treatment effect $\bar{\tau}$ is correct, then $\alpha = 1$. The parameter β measures if heterogeneity in treatment effects is adequately captured. Optimally, $\beta = 1$. If $\beta < 1$, there is overfitting by the causal forest, as predicted deviations of absence spell durations from

[12] Out-of-bag prediction entails using only those trees in the causal or prediction forest where the particular worker was not sampled.

the mean are larger than the actual deviations. If $\beta > 1$, the forest does not capture all of the heterogeneity present. For the forest here, the best linear predictor test results in $\alpha = 1.00$ ($SE = 0.03$) and $\beta = 1.33$ ($SE = 0.07$). Both α and β are close to one, indicating that the estimates adequately capture the average effect of reduced monitoring, as well as deviations from this average in different worker groups. The null hypothesis of no heterogeneity (i.e. $\beta = 0$) is strongly rejected at conventional levels of significance.

Figure 5. *Estimated treatment effects for workers within each quartile of predicted causal forest $\hat{\tau}_x$ estimates*

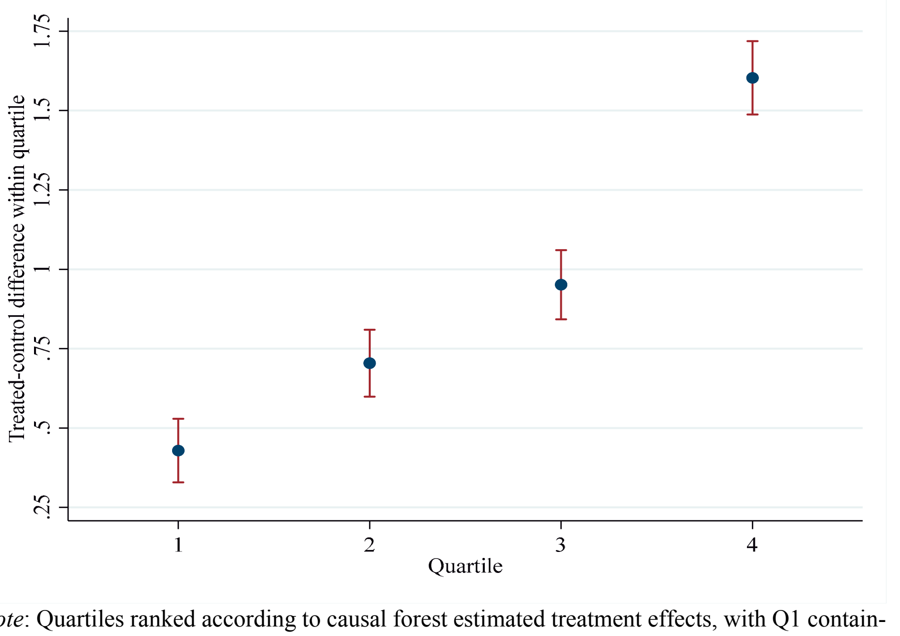

Note: Quartiles ranked according to causal forest estimated treatment effects, with Q1 containing those estimated to be least affected and Q4 those estimated to be most affected. Treatment effects within each of the quartiles estimated as $\hat{\tau} = \bar{y}_i|(W_w = 1) - \bar{y}_i|(W_w = 0)$. Confidence intervals at the 95 percent level shown.

A further test of the causal forest predictions is provided by splitting the data into groups based on the predicted $\hat{\tau}_x$ and then estimating the basic model $\hat{\tau} = \bar{y}_i|(W_w = 1) - \bar{y}_i|(W_w = 0)$ within each group. I subdivide the workers into four quartiles, with Quartile 1 containing those with the smallest predicted treatment effects and Quartile 4 containing those with the largest predicted effects. Estimates of $\hat{\tau}$ within each quartile, along with their 95 percent confidence intervals, are plotted in Figure 5. As expected, treatment effects increase when moving up the quartiles; in addition, treatment effects in all the subgroups are distinct at the 95 percent confidence level. The evidence in Figure 5 points to the effects being fairly small for the majority of the population, but with the group of workers in Quartile 4 increasing their absence spells by 1.5 days or more when monitoring is reduced.

Figure 6. *Survival graphs for absence spells among workers in the held-out test set, ranked by quartiles of predicted treatment effects*

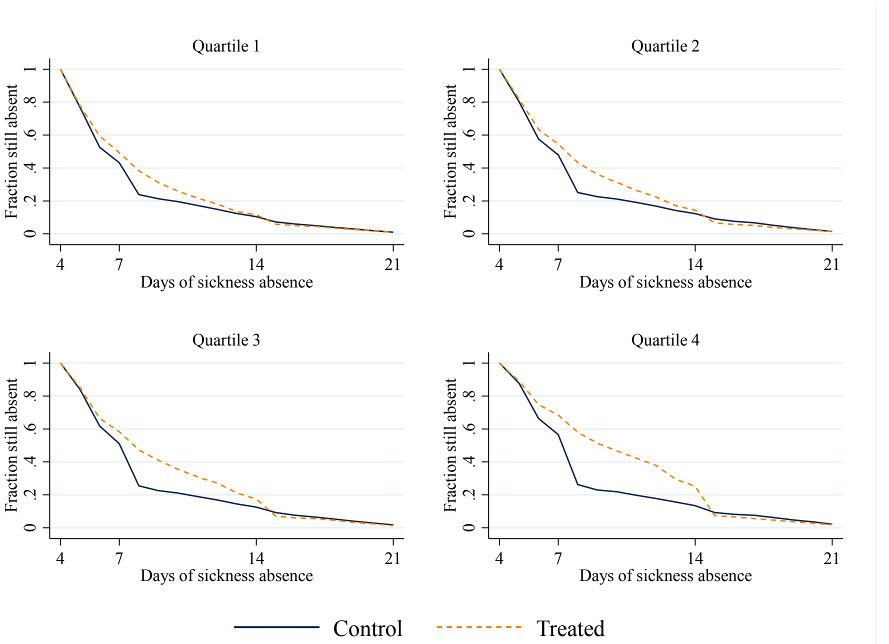

Note: Survival rates for absence spells of the 20 percent of workers randomised into the held-out test set. Workers divided into quartiles based on out-of-bag causal forest predictions. Quartiles ranked according to size of predicted effect, with Q1 containing those estimated to be least affected and Q4 those estimated to be most affected.

As a final check, I use the causal forest to predict treatment effects for workers in the held-out test set and divide them into quartiles based on these predictions, in the same way as was done for workers in the training set. Survival and hazard rates by predicted treatment effect quartile for workers in the test set are presented in Figures 6 and 7 respectively. The graphs confirm that the causal forest has been able to identify workers with different responsiveness to monitoring. The survival curves for treated and control workers in Quartile 1 align fairly closely. A gap between the two groups appears immediately after day 7, but, crucially, it closes completely already before day 14. The maximum difference in the shares of treated and controls still absent from work, on day 8, is 0.12. The gap between treated and controls becomes wider when moving up the predicted treatment effect quartiles. For workers in Quartile 4, a gap opens up already after day 5, and does not close until day 15. The maximum difference between shares of treated and controls absent from work is 0.29, on day 8.

Figure 7. *Hazard graphs for absence spells among workers in the held-out test set, ranked by quartiles of predicted treatment effects.*

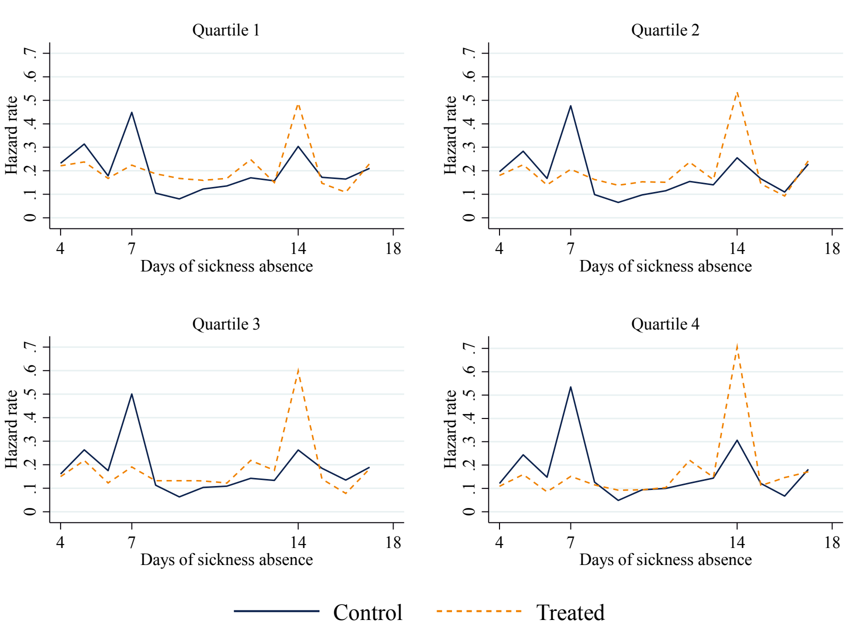

Note: Hazard rates for absence spells of the 20 percent of workers randomised into the held-out test set. Workers divided into quartiles based on out-of-bag causal forest predictions. Quartiles ranked according to size of predicted effect, with Q1 containing those estimated to be least affected and Q4 those estimated to be most affected. The hazard rate represents the probability that a worker who has been absent for a given number of days returns to work on the next day.

The large spikes in hazard rates on days 7 and 14 for workers in the higher quartiles, as shown in Figure 7, are further evidence of behaviour indeed being driven by sensitivity to monitoring. Treated workers in Quartile 1 have a 45 percent hazard rate on day 7, while control workers have a 48 percent hazard rate on day 14. For workers in Quartile 4, the corresponding rates are higher, with 53 percent of controls who are absent on day 7 going back to work on day 8 and 73 percent of treated who are absent on day 14 going back to work on day 15.

6.2 Heterogeneity Drivers

While the ability to predict which workers are likely to change their behaviour when monitoring is relaxed is of high policy relevance, it is also necessary to know which individual characteristics drive these predictions. This is especially so because policymakers are unlikely to have information about all the characteristics used for training the causal forest in this paper, as well as some characteristics being unfeasible to use for moral or legal reasons. One intuitive way of measuring characteristic importance is by assessing how often the trees in the causal forest split on a characteristic. More important characteristics are also more likely to be selected in earlier stages of the tree-growing process.

Such an importance measure is presented in Figure 8. The bars represent the share of splits on each characteristic among splits made by the causal forests' trees, with only the first four rounds of splitting being considered (maximum depth of 4). Each split is weighted by the depth at which it is made, meaning that a split at depth d is given 0.5 the weight of one at $d-1$. According to this metric, the most important variable for treatment effect heterogeneity is the number of days of sick leave in the previous 1.5 years. Other important characteristics are gender, education level, size of establishment and annual earnings. The causal forest almost never split on a number of other variables, such as working in the primary industry, receiving education in teaching, working in the education industry and residing in Gothenburg. The lack of importance assigned to the Gothenburg dummy is interesting, as it is suggestive of the experimental effects working in a similar manner irrespective of whether the rules were tightened or relaxed. Many of the features which have to do with peer effects are assigned some importance, which suggests that there might be some imitation of colleagues' and neighbours' behaviour and attitudes. Results corresponding to Figure 8 when the outcome is the probability of returning to work on days 8-14 are presented in Figure A2 in the Appendix. Qualitatively, the same characteristics are identified as important; however the importance of gender is almost as great as that of sick leave for the probability outcome.

Figure 8. *Importance of worker characteristics for heterogeneity, based on the number of times the causal forest's trees split on the characteristic*

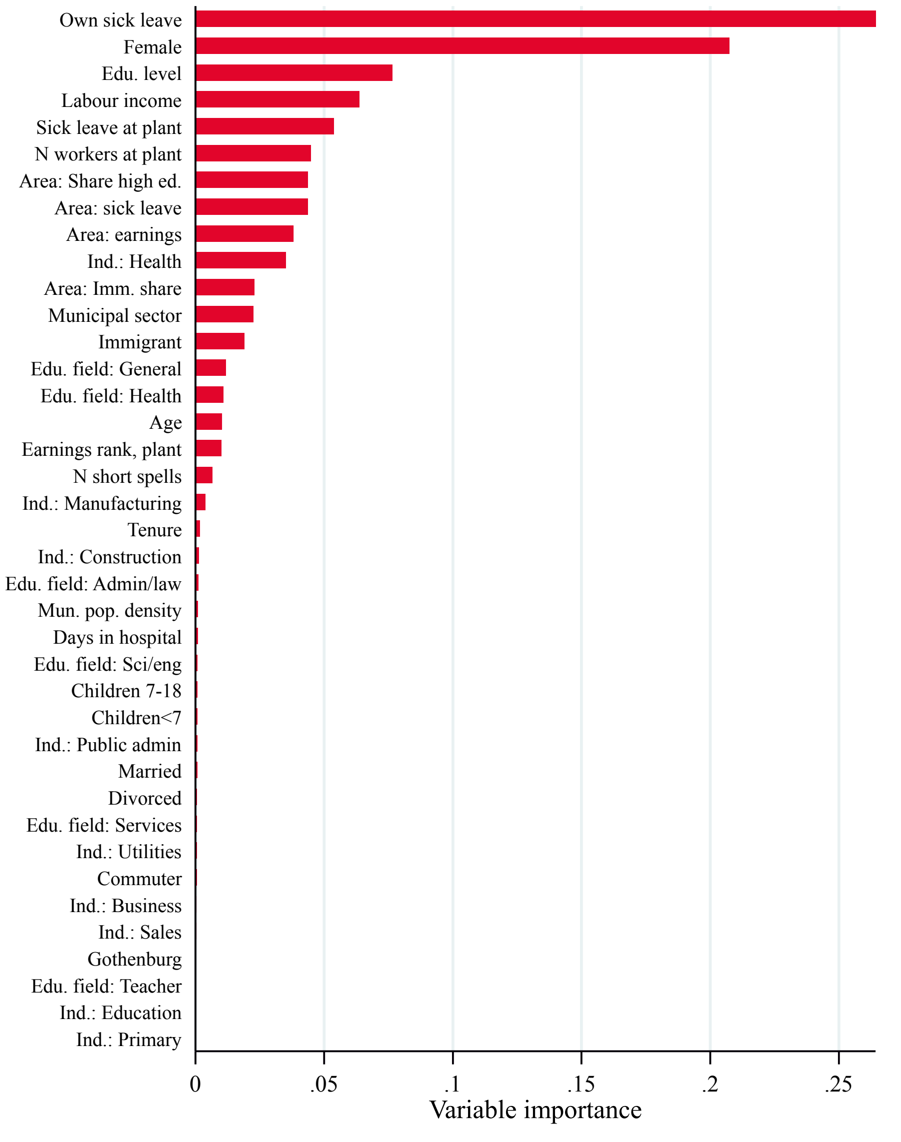

Note: Importance is measured as share of splits at maximum depth of 4 within the trees. Splits at lower depth d given two times the weight of those at $d+1$. Total importance sums to 1.

Another way of measuring the contribution of different characteristics is by evaluating the partial dependence function, that is the mean of causal forest predictions when the value of one of the x_k is changed to some $\bar{\bar{x}}$ over the distribution of the other $\boldsymbol{x}_{-k}$. The partial dependence function $\hat{f}_{\boldsymbol{x}_{-k},\ x_k=\bar{\bar{x}}} =$

$E_{\boldsymbol{x}_{-k}}(\hat{\tau}_{\boldsymbol{x}_{-k},\ x_k=\bar{\bar{x}}}) = \int \hat{\tau}_{\boldsymbol{x}_{-k},\ x_k=\bar{\bar{x}}}\, \mathrm{d}\mathbb{F}_{\boldsymbol{x}_{-k}}$ is estimated on the training dataset as:

$$\hat{f}_{\boldsymbol{x}_{-k},\ x_k=\bar{\bar{x}}} = \frac{\sum_{i=1}^{N} \hat{\tau}_{\boldsymbol{x}_{-k}=\boldsymbol{x}_{-k,w}\ ,\ x_k=\bar{\bar{x}}}}{N} \tag{7}$$

$\boldsymbol{x}_{-k} = \boldsymbol{x}_{-k,w}$ indicates that the other characteristics are held at their true values for the worker w. The empirically observed distribution of the $\boldsymbol{x}_{-k}$ thus serves as an approximation of the population distribution.

Partial dependence plots for all included worker characteristics are provided in the Appendix. For the variables which are continuous and take on many values in the dataset, $\hat{f}_{\boldsymbol{x}_{-k},\ x_k=\bar{\bar{x}}}$ is estimated at $\bar{\bar{x}}$ equal to their values at each decile and is plotted in Figure A3. For binary covariates, $\hat{f}_{\boldsymbol{x}_{-k},\ x_k=\bar{\bar{x}}}$ is evaluated at the values 0 and 1, and is plotted in Figure A4. As only 6.7 percent of workers have spent any days in hospital during the pre-period, the effect of that characteristic is evaluated at 0 and 9 days (the mean among those with any hospitalisations) and is shown together with the binary characteristics in Figure A4. Partial dependence plots for ordinal or discrete covariates, evaluated at each of their levels, are presented in Figure A5. The figure also contains the plot for municipal population density (ordered from lowest to highest among the nine municipalities included in the experiment).

The plots in Figures A3-A5 are reassuring as causal forest estimates vary more with variables on which the causal forest has often split. The highest and lowest mean predictions, 0.77 and 1.16, are obtained when the days of previous sick leave are set to their 1st and 9th decile levels of 2 and 117, respectively. Other variables which the forest has often used for splitting, such as gender, education level, annual earnings, as well as number of workers and average sick leave at the worker's establishment also affect $\hat{f}_{\boldsymbol{x}_{-k},\ x_k=\bar{\bar{x}}}$.

As both variable importance metrics draw attention to the sick leave history variable, I plot the full distribution of causal forest estimates of $\hat{\tau}_x$ when it is reassigned in Figure 9. Days of sick leave in previous periods are set to the values 2 and 117, the levels at the 1st and 9th decile levels respectively. If all workers behaved as if they had high pre-period sick leave, the average $\hat{\tau}_x$ would increase by 0.22 days to 1.16, as shown in the left panel. The entire distribution of $\hat{\tau}_x$ would shift to the right. The opposite holds when all workers are re-assigned to have low pre-period sick leave in the right panel; the average $\hat{\tau}_x$ falls by 0.17 days and the entire upper tail of the distribution disappears, as practically no workers are estimated to have $\hat{\tau}_x$ in excess of 1.5. Importantly, even if all workers had low pre-period sick leave, almost none of them are estimated to reduce duration of sickness absence spells when monitoring is reduced. This showcases a key advantage of the causal forest relative to linear models, which are likely to make predictions that are theoretically unreasonable when parameters are changed.

It is important to note that setting a characteristic x_k to a value $\bar{\bar{x}}$ while the other variables are held at their empirical values might be problematic. Some of the $\boldsymbol{x}_{-k}$ might covary in a natural way with x_k. For example, neighbourhoods with high average earnings also tend to have high shares of highly educated individuals. For this reason, I vary groups of interrelated variables at the same time and evaluate how the entire distribution of predicted τ_x changes in response. This exercise is guided by the importance results presented in Figure 8 and Figures A4-A5. Four groups of variables are considered: education and annual earnings; gender; neighbourhood earnings, immigrant share, share with post-secondary education and average days of sick leave; and workplace size and average sickness absence in previous periods. The results are shown in Figure 10.

Figure 9. *Treatment effect estimates when all workers are assigned to have 117 days of previous sick leave (90th percentile) and 2 days of previous sick leave (10th percentile)*

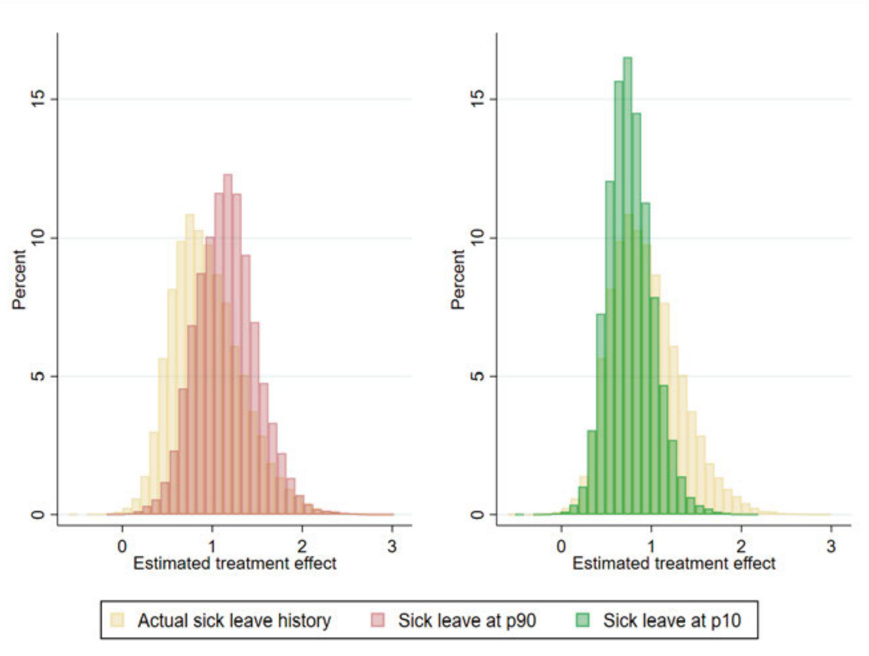

Note: Bars represent percentage of spells falling into each treatment effect bin. All worker characteristics except sick leave history held at their empirical values.

The upper left hand panel presents how $\hat{\tau}_x$ estimates change when socioeconomic status is varied. The green line shows the distribution of estimated effects if all workers behaved as if they had earnings of 145,953 SEK (the 90th percentile level) and some college education. Only 10 percent of workers are estimated to have $\hat{\tau}_x$ in excess of 1.27 days, compared to 19 percent of the $\hat{\tau}_x$ if workers are assigned their actual earnings and education. The distribution of predicted treatment effects if all workers behaved as if they had annual earnings of 44,311 SEK (the 10th percentile level), as well as less than nine years of education, is given by the red line. Only 10 percent of treatment effects are smaller than 0.78 days, compared to 38 percent if workers have their observed earnings and education.

Figure 10. *Cumulative distribution of treatment effect estimates if workers would behave as if they had different socioeconomic status, gender, neighbourhood characteristics and workplace characteristics.*

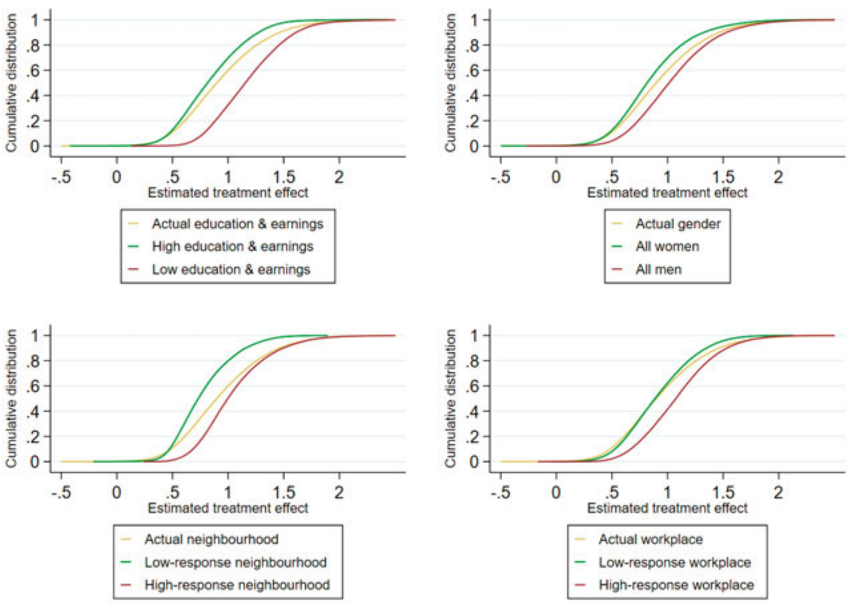

Note: *Education & earnings*: All workers assigned to have annual earnings of 145,953 SEK and some college education (high socioeconomic status); annual earnings of 44,311 SEK and less than compulsory education (low socioeconomic status). *Gender*: All workers assigned as men and women. *Neighbourhood*: Average earnings at 152,113 SEK, share highly educated at 38.6 percent, average sick days 11.1, immigrant share 3.7 percent (low-response neighbourhoods); Average earnings at 106,901 SEK, share highly educated at 4.6 percent, average sick days 33.4, immigrant share 43.8 percent (high-response neighbourhoods). *Workplace*: Number of workers at workplace 11, days of sickness absence at workplace 5.5 (low-response workplaces); number of workers at workplace 10726, days of sickness absence at workplace 18.1 (high-response workplaces).

The green and red lines in the top right panel of Figure 10 show the cumulative distribution function of predicted $\hat{\tau}_x$ if all the workers in the sample had behaved as if they were women and men respectively. The differences are sizeable, but somewhat smaller than for the other variable groups. Larger effects can be seen when neighbourhood variables are varied in the bottom left panel. The green line shows the distribution if neighbourhood earnings and share of highly educated are set to their 90th percentile values (152,113 SEK and 38.6 percent) and average days of sick leave and immigrant share are set to their 10th percentile values (11.1 and 3.7 percent). The red line illustrates the case when earnings and share highly educated are set to their 10th percentile values (106,901 SEK and 4.6 percent) and average sick leave and immigrant share to their 90th percentile values (33.4 and 43.8 percent). Effects are sizeable, especially when all workers are re-assigned to live in neighbourhoods high on the socioeconomic scale. Finally, the bottom right panel plots the cumulative distribution function when establishment size is set to be small (11 workers) and establishment sick leave to be low (5.5 days/worker) in comparison to when the establishment size is set to be large (10,726 workers) and sick leave to be high (18.1 days/worker). Effects are smaller than when the other variables are varied, but the $\hat{\tau}_x$ distribution clearly shifts to the right if workers are re-coded to be at large establishments with high sick leave.

6.3 Characterising Workers Who Are Sensitive to Monitoring

To characterise workers who are sensitive and non-sensitive more fully, I once again divide the sample into four quartiles, where Quartile 1 contains those with the smallest $\hat{\tau}_x$ and Quartile 4 contains those with the largest $\hat{\tau}_x$. Average values of each characteristic among workers in each of the quartiles are shown Table 1, with the ordering of characteristics guided by the importance measure in Figure 8. Variables identified as important vary in a systematic way across quartiles, suggesting that key relationships in the data have been correctly identified. Several points stand out. Firstly, workers in the more affected quartiles have taken out substantially more sick leave in previous periods than workers in the less affected quartiles. On average, workers in Quartile 4 have taken out 87 days of sick leave over the previous 1.5 years, compared to 23 days for workers in Quartile 1. Secondly, the most responsive workers are characterised by having low socioeconomic status. They are less educated, have lower earnings and are more likely to be immigrants than workers in the other quartiles. This pattern is also found in the characteristics of the neighbourhoods where they live, which are characterised by low earnings and education and high immigrant shares and sick leave uptake. The correlation between high neighbourhood sick leave uptake and high responsiveness to monitoring suggests that peer effects might come into play, as does the apparent connection with colleagues' sickness absence. While it is well-established that women tend to take out more sick leave than men, and that sickness absence is more prevalent in the public sector, the results provide no evidence that these predictors of sick leave uptake are correlated with responsiveness to monitoring. On the contrary, women constitute only 27 percent of the most responsive quartile, compared to 76 percent of the least responsive one.

Table 1. *Averages of characteristics in each predicted treatment effect quartile.*

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Own sick leave	25.67	35.64	48.67	86.86
	(0.184)	(0.245)	(0.28)	(0.399)
Female	0.783	0.639	0.431	0.217
	(0.002)	(0.003)	(0.003)	(0.002)
Edu. level	3.694	3.022	2.633	2.366
	(0.008)	(0.007)	(0.007)	(0.006)
Labour income	108900	97100	97250	86910
	(300.7)	(235.6)	(221.1)	(224.9)
Sick leave at plant	10.66	11.56	12.71	14.41
	(0.028)	(0.032)	(0.034)	(0.037)
N workers at plant	1689	1801	2201	3384
	(19.04)	(21.39)	(23.41)	(28.18)
Area: sick leave	16.29	19.75	22.65	24.5
	(0.037)	(0.046)	(0.05)	(0.051)
Area: earnings	136500	126800	120900	118000
	(124.7)	(105.7)	(92.48)	(74.21)
Area: Share high ed.	0.259	0.192	0.149	0.128
	(0.001)	(0.001)	(0.001)	(0.001)
Ind.: Health	0.507	0.383	0.249	0.131
	(0.003)	(0.003)	(0.002)	(0.002)
Municipal sector	0.583	0.469	0.34	0.244
	(0.003)	(0.003)	(0.003)	(0.002)
Area: Imm. share	0.13	0.174	0.221	0.252
	(0.001)	(0.001)	(0.001)	(0.001)
Immigrant	0.103	0.227	0.361	0.487
	(0.003)	(0.004)	(0.005)	(0.005)
Edu. field: Health	0.254	0.128	0.057	0.026
	(0.002)	(0.002)	(0.001)	(0.001)
Age	34.82	34.9	35.71	36.62
	(0.066)	(0.069)	(0.069)	(0.068)
Earnings rank, plant	0.596	0.559	0.542	0.484
	(0.001)	(0.001)	(0.001)	(0.001)
Edu. field: General	0.27	0.431	0.513	0.575
	(0.003)	(0.003)	(0.003)	(0.003)
N short spells	4.179	4.393	4.96	5.864
	(0.021)	(0.021)	(0.021)	(0.025)
Ind.: Manufacturing	0.131	0.194	0.305	0.439
	(0.002)	(0.002)	(0.003)	(0.003)
Tenure	1.45	1.442	1.497	1.502
	(0.007)	(0.007)	(0.007)	(0.007)

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Ind.: Construction	0.022	0.046	0.082	0.097
	(0.001)	(0.001)	(0.002)	(0.002)
Mun. pop. density	802.1	784.8	803.8	825.3
	(2.034)	(2.119)	(2.029)	(1.915)
Edu. field: Sci/eng	0.143	0.191	0.266	0.295
	(0.002)	(0.002)	(0.003)	(0.003)
Edu. field: Admin/law	0.211	0.155	0.101	0.061
	(0.002)	(0.002)	(0.002)	(0.001)
Days in hospital	0.362	0.516	0.583	1.082
	(0.016)	(0.026)	(0.024)	(0.051)
Children<7	0.224	0.223	0.221	0.188
	(0.003)	(0.003)	(0.003)	(0.003)
Children 7-18	0.405	0.418	0.389	0.348
	(0.004)	(0.004)	(0.004)	(0.004)
Married	0.496	0.473	0.481	0.427
	(0.003)	(0.003)	(0.003)	(0.003)
Ind.: Public admin	0.023	0.019	0.017	0.046
	(0.001)	(0.001)	(0.001)	(0.001)
Ind.: Utilities	0.054	0.071	0.095	0.096
	(0.001)	(0.001)	(0.002)	(0.002)
Edu. field: Services	0.054	0.059	0.049	0.038
	(0.001)	(0.001)	(0.001)	(0.001)
Divorced	0.094	0.108	0.118	0.139
	(0.002)	(0.002)	(0.002)	(0.002)
Commuter	0.115	0.101	0.098	0.093
	(0.002)	(0.002)	(0.002)	(0.002)
Ind.: Business	0.1	0.082	0.068	0.047
	(0.002)	(0.002)	(0.001)	(0.001)
Ind.: Sales	0.13	0.167	0.15	0.111
	(0.002)	(0.002)	(0.002)	(0.002)
Gothenburg	0.83	0.812	0.832	0.855
	(0.002)	(0.002)	(0.002)	(0.002)
Edu. field: Teacher	0.068	0.037	0.015	0.006
	(0.001)	(0.001)	(0.001)	(0)
Ind.: Education	0.027	0.028	0.022	0.016
	(0.001)	(0.001)	(0.001)	(0.001)
Ind.: Primary	0.003	0.006	0.006	0.009
	(0)	(0)	(0)	(0.001)

Note: Quartiles ranked according to causal forest estimated treatment effects, with Q1 containing those estimated to be least affected and Q4 those estimated to be most affected. Colours are assigned according to how strongly the average value of the variable in the quartile deviates from its grand mean.

The corresponding figures for municipal employees are 24 and 63 percent. A corollary of this result is the high share of manufacturing workers, 48 percent, in Quartile 4. On the other hand, two other well-known predictors of sickness absence, establishment size and the worker's low position within the establishment,[13] are positively correlated with strong responses to monitoring.

Defining the four quartiles based on the probability of returning to work on days 8-14 has little effect on which groups of workers are identified as sensitive to monitoring, as shown in Table A2 in the Appendix. Individuals with high pre-period sick leave uptake, men, those with low earnings and education, and those who live in socioeconomically disadvantaged neighbourhoods are once again overrepresented in Quartile 4.

The high predictive power of the individual's sick leave history on responsiveness to monitoring is encouraging, as this characteristic is readily available to policymakers and its use is not counter to legal restrictions. Two other characteristics that are of interest for policymakers in directing monitoring efforts are plant size and the average sickness absence at a plant. The relevance of these variables also indicates that there is a case to be made for inducing employers with high sick leave uptake among the workforce to provide more healthy working environments.

6.4 Targeted Monitoring Policy

The causal forest predictions of $\hat{\tau}_x$ can be used for selective monitoring of workers. Since monitoring is costly, especially because it takes up medical professionals' time, it may be beneficial to monitor workers with high $\hat{\tau}_x$ more and workers with low $\hat{\tau}_x$ less. I assess the gains of such a selective monitoring policy relative to randomising monitoring intensity as in the experiment. The results, based on the held-out test set of workers, are presented in Figure 11.

The aim is to select the policy which leads to the largest reduction in sickness absence for a given share s of workers that is monitored at a higher intensity (i.e. after 7 as opposed to 14 days). The reduction in sickness absence a is in terms of the share of the total reduction that is achieved if all workers are monitored more intensely; monitoring all workers would reduce sickness absence by $a = 1$. The blue line shows the reduction in sickness absence duration if the workers in s are randomly selected. It oscillates around the 45° line; in expectation, absence duration is reduced by a share $a = s$. Using the $\hat{\tau}_x$ allows for greater efficiency. If workers are ranked according to $\hat{\tau}_x$ and s contains the $s \times N$ workers with the highest $\hat{\tau}_x$, the reduction in absence for given s is indicated by the red line. This lies above the blue line, showing that it is possible to reduce sickness absence by $a > s$. Finally, the green line is the corresponding performance that can be achieved if the policymaker only

[13] Proxied for by the worker's earnings rank within the establishment.

has information about previous days of sickness absence. In this policy, workers are sorted into bins based on the number of days of sickness absence they took out in the past. Workers with higher past sickness absence are monitored first.[14] Most of the gains of the full causal forest model are retained if only sick leave history information is used. In particular, consider the case of monitoring the same share of workers as in the experiment, $s = 0.51$. In expectation, randomising who gets monitored gives $a = 0.51$.

Figure 11. *Effect in terms of reducible sickness absence duration for given share of workers monitored according to different monitoring policies*

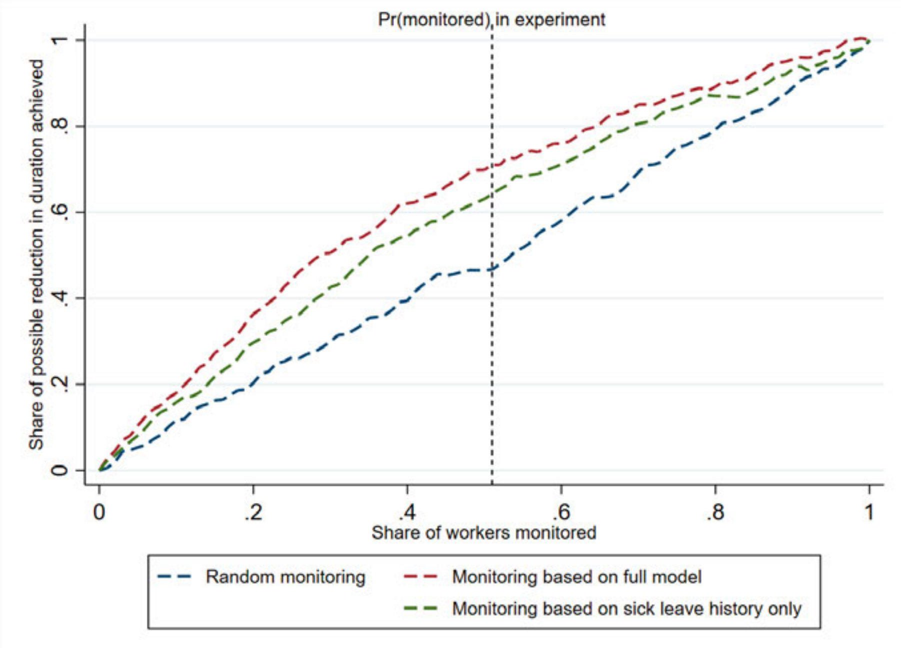

Note: Effects of monitoring test sample workers according to different rules. Monitoring based on full model assumes workers are ranked based on their estimated treatment effects and those with higher estimated treatment effects are monitored first. Monitoring based on sick leave history only assumes workers are ranked based on the estimated treatment effect of their sick leave history bin. Workers in bins with higher estimated treatment effects are monitored first. Order of monitoring within a sick leave history bin is random. Pr(monitored) in experiment = 0.51.

However, monitoring the test set workers using the other policies can achieve significantly larger decreases in sickness absence. The full-information policy is estimated to decrease the reducible part of sickness absence by $a = 0.71$ when $s = 0.51$; the policy which only uses sick leave history yields $a = 0.65$ for $s = 0.51$. The full-information policy thus allows reducing monitoring for the same share of workers as in the experiment for a 41 percent smaller loss

[14] Selection of which workers are monitored first within a bin is random.

in terms of extra sickness absence; the sick leave history policy still results in a 29 percent smaller loss. Assuming monitoring costs are constant across workers, the targeted policies can save the same amount of resources without increasing sickness absence nearly as much.

7. Conclusion

There is a strong case for ensuring that the sickness insurance system is fair, adequately compensating those who have temporarily lost the ability to work, while providing minimal incentives for overuse by healthy individuals. A common way of attaining this goal is by having qualified medical professionals monitor recipients. However, the opportunity cost of these professionals' time is high and it is important to know where it is put to the best use. This paper attempts to shed light on this question by assessing which workers' behaviour responded the most when medical certificate requirements were relaxed in a randomised experiment.

The evidence points to substantial heterogeneity in worker responses. Sickness absence spells are estimated to have been only 0.36 days longer for the least sensitive decile of individuals, compared to 1.71 days longer for the most sensitive decile. The key predictors of strong behavioural changes when monitoring intensity is varied are high previous sick leave uptake, low socioeconomic status and male gender. There is also evidence that colleagues' and neighbours' behaviour has an effect. A key finding is that many predictors of high sick leave uptake, such as female gender and working in the public sector, are not predictors of high sensitivity to monitoring. The existence of workplaces with high sick leave uptake and high monitoring responsiveness suggests that the management at such establishments should take steps to improve working conditions. This is especially pertinent in light of findings that such measures are effective in reducing absenteeism (Huber et al., 2015).

For policymakers, selective monitoring of sickness insurance recipients can be a way of reducing costs while minimising the effect on sickness absence uptake. Back-of-the-envelope calculations suggest that monitoring could be reduced by the same amount as in the experiment, but causing only 59 percent of the increase in sickness absence if efforts are targeted using all the characteristics included in this study, or 71 percent of the increase if only sick leave history is used.

While targeted monitoring has high potential when it comes to increasing efficiency, ethical concerns must also be taken into account when designing policy. Monitoring based on many of the worker characteristics included in the full causal forest model would likely be seen as discriminatory or unfair. In particular, it would be highly controversial to use variables such as gender, immigrant background, or income for monitoring purposes. A policy which varies monitoring intensity based only on sick leave history would thus be preferable for ethical and practical reasons. Another upside of such a policy is that it self-regulates against overuse by individuals who have little past sick leave uptake. If these workers increase their sickness absence by a significant amount in response to the reduction in monitoring, they will eventually end up in the more intensely monitored group.

Another concern to keep in mind is that not all reductions in sick leave are socially beneficial. If there are monetary or time costs of obtaining a medical certificate, workers might forgo days of absence which would have been medically motivated. This might lead to both negative longer-term effects on the worker's own health (see e.g. Marie and Vall Castelló, forthcoming) and to the infection of others at the workplace, an issue which has been prominent during the Covid-19 pandemic. The design of a policy which takes these broader issues into account is left for future research.

References

Angelov, Nikolay, Johansson, Per, Lindahl, Erica, Lindström, Elly-Ann. (2011). Kvinnors och mäns sjukfrånvaro. *IFAU Report* 2011:2.

Athey, Susan, and Imbens, Guido. (2016). Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113.27:7353-7360.

Athey, Susan, Tibshirani, Julie, Wager, Stefan. (2019). Generalized random forests. *Ann. Statist*. 47.2:1148 – 1178.

Barmby, Tim A., Ercolani, Marco G., Treble, John G.. (2002). Sickness Absence: An International Comparison. *The Economic Journal*, 112:480.

Bedard, Kelly, and Dhuey, Elizabeth. (2006). The persistence of early childhood maturity: International evidence of long-run age effects. *The Quarterly Journal of Economics* 121.4: 1437-1472.

Boeri, Tito, Di Porto, Edoardo, Naticchioni, Paolo, Scrutinio, Vincenzo. (2021). Friday Morning Fever. Evidence from a Randomized Experiment on Sick Leave Monitoring in the Public Sector. *CEPR Discussion Paper* No. DP16104.

Bratberg, Espen, and Monstad, Karin. (2015). Worried sick? Worker responses to a financial shock. *Labour Economics*, 33:111-120

Breiman, Leo. (2001). Random forests. *Mach. Learn*. 45:5–32.

Breiman, Leo, Friedman, Jerome H., Olshen, Richard. A. and Stone, C. J. (1984). *Classification and Regression Trees*. Wadsworth Advanced Books and Software, Belmont, CA.

Chernozhukov, Victor, Demirer, Mert, Duflo, Esther and Fernàndez-Val, Ivàn. (2020). Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, with an Application to Immunization in India. *National Bureau of Economic Research*, WP 24678.

Ferman, Bruno, Torsvik, Gaute, Vaage, Kjell. (2021). Skipping the doctor: evidence from a case with extended self-certification of paid sick leave. *J Popul Econ* 1-37.

Försäkringskassan. (2021). *Effekter som covid-19 har på sjukförsäkringen*, FK 2020/000065.

Frick, Bernd, and Malo, Miguel Á. (2008). Labor market institutions and individual absenteeism in the European Union: the relative importance of sickness benefit systems and employment protection legislation. *Industrial Relations: A Journal of Economy and Society*, 47.4:505-529.

Hartman, Laura, Hesselius, Patrik, Johansson, Per. (2013). Effects of eligibility screening in the sickness insurance: Evidence from a field experiment, *Labour Economics*, 20:48-56

Henrekson, Magnus, and Persson, Mats. (2004). The effects on sick leave of changes in the sickness insurance system. *Journal of Labor economics* 22.1:87-113.

Hensvik, Lena and Rosenqvist, Olof. (2019). Keeping the Production Line Running: Internal Substitution and Employee Absence, *J. Human Resources* 54:200-224

Hesselius, Patrik, Nilsson, Peter J., Johansson, Per. (2009) Sick of Your Colleagues' Absence?, *Journal of the European Economic Association*, 7.2-3.

Hesselius, Patrik, Johansson, Per, Vikström, Johan. (2013). Social behaviour in work absence. *The Scandinavian Journal of Economics,* 115.4:995-1019.

Huber, Martin, Lechner, Michael, Wunsch, Conny. (2015). Workplace health promotion and labour market performance of employees. *Journal of Health Economics*, 43:170-189.

Johansson, Per, Karimi, Arizo, Nilsson, Peter J. (2019). Worker absenteeism: peer influences, monitoring and job flexibility. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 182.2:605-621.
Johansson, Per, and Palme, Mårten. (2002). Assessing the effect of public policy on worker absenteeism. *Journal of Human Resources*, 37.2:381-409.
Johansson, Per and Palme, Mårten. (2005). Moral hazard and sickness insurance. *Journal of Public Economics*, 89.9–10:1879-1890.
Knaus, Michael C., Lechner, Michael, and Strittmatter, Anthony. (2021). Machine learning estimation of heterogeneous causal effects: Empirical Monte Carlo evidence. *The Econometrics Journal* 24.1:134-161.
Lindbeck, Assar, Palme, Mårten, and Persson, Mats. (2016). Sickness absence and local benefit cultures. *The Scandinavian Journal of Economics*, 118.1:49-78.
Lindgren, Karl-Oskar, *Workplace size and sickness absence transitions*, IFAU Working paper 2012:26.
Marie, Olivier and Vall Castelló, Judit. (forthcoming). Sick Leave Cuts and (Unhealthy) Returns to Work. *Journal of Labor Economics*.
Nie, Xinkun, and Wager, Stefan. (2021). Quasi-oracle estimation of heterogeneous treatment effects. *Biometrika*, 108.2:299-319.
van Ommeren, Jos N. and Gutiérrez-i-Puigarnau, Eva. (2011). Are workers with a long commute less productive? An empirical analysis of absenteeism, *Regional Science and Urban Economics*, 41.1.
OECD. (2020). *Paid sick leave to protect income, health and jobs through the COVID-19 crisis*.
OECD. (2021). *Public spending on incapacity (indicator)*.
Palme, Mårten, and Persson, Mats. (2020). Sick Pay Insurance and Sickness Absence: Some European Cross- Country Observations and a Review of Previous Research. *Journal of Economic Surveys* 34.1:85-108.
Riksförsäkringsverket. (1989). *Utvidgad egen sjukskrivning*, AD 1989-954:01.
SOU 1981:22, *Sjukförsäkringsfrågor: Betänkande av sjukpenningkommitén*.
SOU 1991:68, *Frikommunförsöket: Erfarenheter av försöksverksamheten med avsteg från statlig reglering m.m*.
SOU 2015:21, *Mer trygghet och bättre försäkring*.
Treble, John and Barmby, Tim. (2011). *Worker absenteeism and sick pay*, Cambridge University Press, Cambridge/New York.
Winkelmann, Rainer. (1999). Wages, firm size and absenteeism, *Applied Economics Letters*, 6:6, 337-341.

Appendix

Table A1. *Balancing table for treated and control worker characteristics*

	Controls (N=99,103)			Treated (N=95,276)			
	Mean	**Min**	**Max**	**Mean**	**Min**	**Max**	**Diff.**
Age	37.634	18	64	37.578	18	64	0.056
Female	0.498	0	1	0.496	0	1	0.002
Native	0.870	0	1	0.874	0	1	-0.004**
Immigrant:							
Nordic	0.053	0	1	0.053	0	1	0.000
Other Europe	0.049	0	1	0.048	0	1	0.001
Rest of World	0.028	0	1	0.025	0	1	0.003***
Married	0.531	0	1	0.531	0	1	0.000
Divorced	0.095	0	1	0.094	0	1	0.000
Children younger than 7	0.222	0	5	0.222	0	5	-0.000
Children aged 7-18	0.453	0	7	0.451	0	7	0.002
Population density in municipality	764.406	1.088	963.289	764.959	1.088	963.289	-0.553
Gothenburg	0.791	0	1	0.791	0	1	-0.001
Neighbourhood:							
Average annual earnings	130112	33600	933600	130303	25500	933600	-191.76*
Share with post-secondary education	0.209	0	0.800	0.210	0	0.800	-0.001*
Average days of sick leave in previous 1.5 years	18.588	0	318	18.521	0	185	0.067*
Immigrant share	0.160	0	1	0.159	0	1	0.001
Education level	3.206	0	7	3.207	0	7	-0.001
Education field:							
General	0.390	0	1	0.388	0	1	0.001
Teacher	0.032	0	1	0.032	0	1	-0.000
Administration, law, social science	0.158	0	1	0.162	0	1	-0.004**
Science and engineering	0.224	0	1	0.224	0	1	0.000
Health	0.128	0	1	0.127	0	1	0.002
Services	0.045	0	1	0.045	0	1	0.000
Annual labour income	109042	22962	2299756	109087	22966	2263504	-44.086
Commuter to another municipality	0.119	0	1	0.123	0	1	-0.004**
Municipal sector	0.395	0	1	0.391	0	1	0.004*

	Controls (N=99,103)			Treated (N=95,276)			
	Mean	**Min**	**Max**	**Mean**	**Min**	**Max**	**Diff.**
Industry:							
Primary	0.009	0	1	0.009	0	1	-0.000
Manufacturing	0.225	0	1	0.224	0	1	0.001
Construction	0.057	0	1	0.058	0	1	-0.001
Utilities	0.073	0	1	0.073	0	1	-0.000
Sales	0.168	0	1	0.166	0	1	0.002
Business services	0.098	0	1	0.100	0	1	-0.002*
Health	0.306	0	1	0.307	0	1	-0.001
Education	0.031	0	1	0.030	0	1	0.001*
Public administration	0.025	0	1	0.024	0	1	0.000
Plant:							
Number of workers at plant	1759.6	1	38420	1766.1	1	38420	-6.471
Earnings rank at plant	0.587	0.009	1	0.587	0.012	1	-0.000
Tenure	1.554	0	3	1.557	0	3	-0.003
Average days of sick leave at plant in past 1.5 years	10.976	0	120.500	10.945	0	176	0.031
Share of plant employed treated	0.490	0	1	0.489	0	1	0.000
Days of sick leave in past 1.5 years	37.903	0	1256	38.281	0	1162	-0.378
Days in hospital in past 1.5 years	0.542	0	376	0.564	0	414	-0.022
Number of short spells in past 1.5 years	3.019	0	58	2.984	0	36	0.036**

Note: Statistics for treated and control workers who fulfil the restrictions on being included in the main analysis, that is are aged 18-64, have annual earnings at least three times a "minimum" monthly wage (defined as the tenth percentile among blue-collar workers) and do not work for the central government. Workers who did not have any sickness absence spells during the experimental period are included.

Figure A1. *The distribution of predicted treatment effects on the probability of a sickness absence spell ending on days 8-14.*

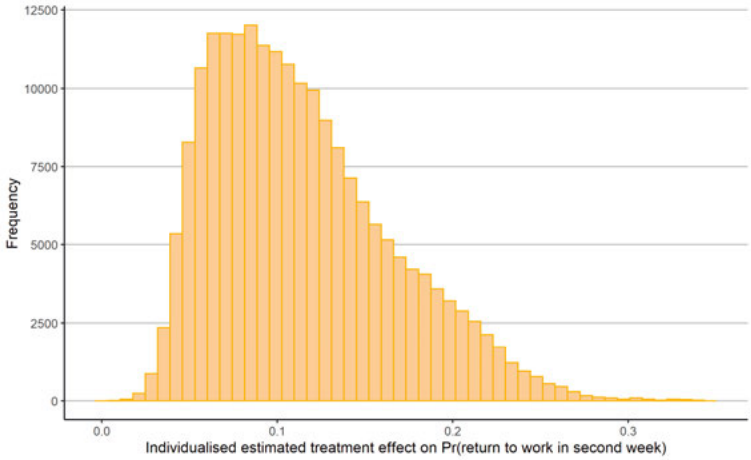

Note: Frequency represents the number of sickness spells with estimated causal forest treatment effects that fall within each bin.

Figure A2. *Importance of each of the characteristics considered for determining heterogeneity in probability of returning to work on days 8-14.*

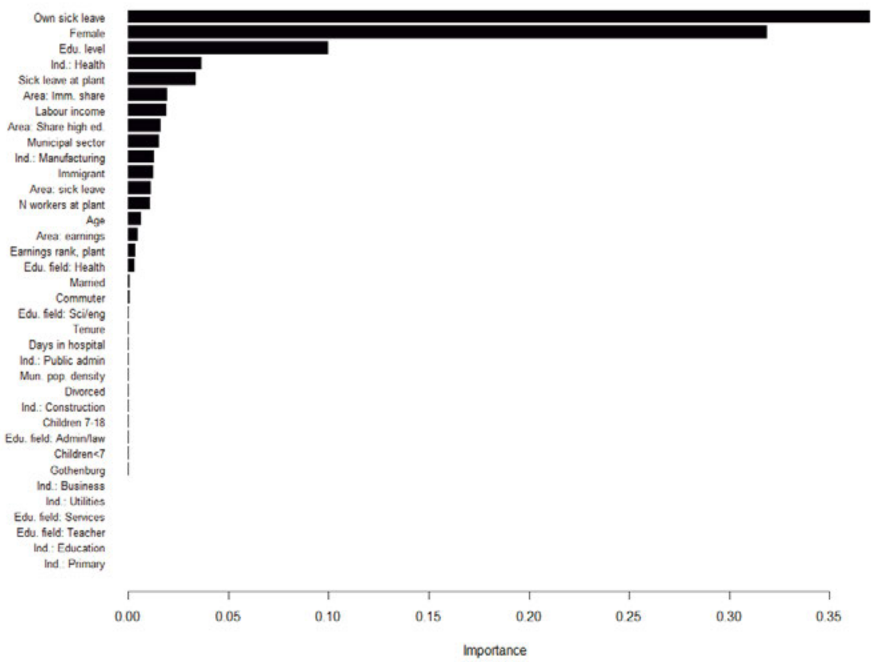

Note: Importance is measured as share of splits at maximum depth of 4 within the trees. Splits at lower depth d given 0.5 the weight of those at $d + 1$. Total importance sums to 1.

Figure A3. *Partial dependence plots for the continuous covariates.*

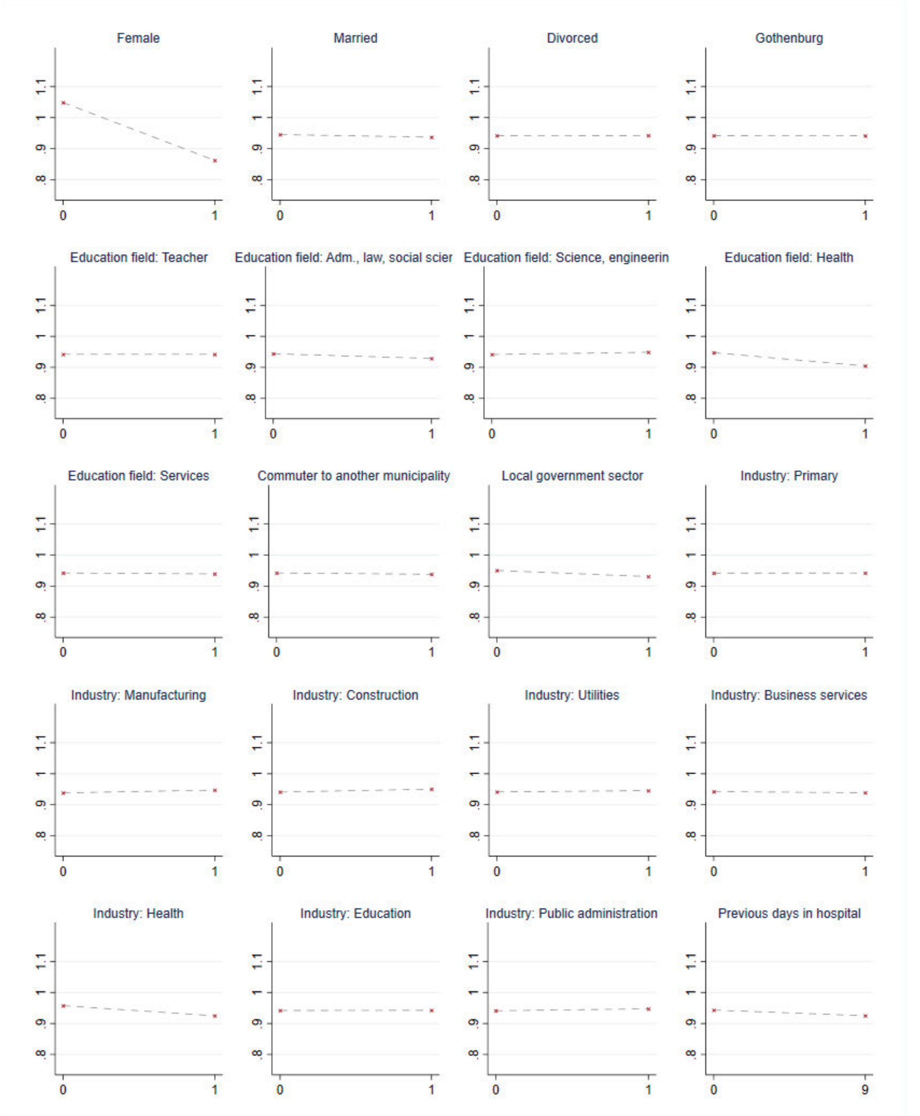

Note: Increase in spell duration in days (*y*-axis) evaluated when the covariate takes on the value 0 or 1 (*x*-axis). For days in hospital, increase in spell duration evaluated at 0 and at 9 (the average number of days among those with any previous hospitalisation). This is because only 6.8 percent of the sample have spent any number of days in hospital.

Figure A4. *Partial dependence plots for the categorical or discrete covariates.*

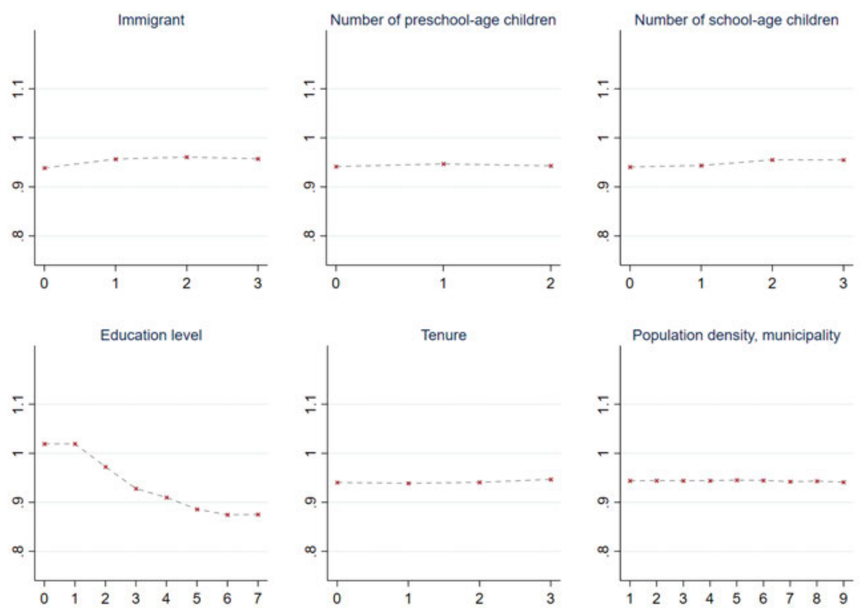

Note: Increase in spell duration in days (*y*-axis) evaluated when the covariate takes on its different possible values (*x*-axis). Individuals with >2 pre-school age children and >3 school age children present in the data, but effects not evaluated due to their small proportion. For population density, each category represents the density in one municipality, ordered from least densely populated to most densely populated.

Table A2. *Averages of characteristics in each predicted treatment effect quartile when the outcome is probability of returning to work on days 8-14.*

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Own sick leave	15.62	34.25	50.88	77.33
	(0.098)	(0.201)	(0.262)	(0.291)
Female	0.823	0.722	0.572	0.149
	(0.002)	(0.002)	(0.002)	(0.002)
Edu. level	4.156	3.258	2.618	2.401
	(0.006)	(0.006)	(0.005)	(0.005)
Ind.: Health	0.57	0.459	0.322	0.101
	(0.002)	(0.002)	(0.002)	(0.001)
Sick leave at plant	10.2	11.2	12.48	14.13
	(0.019)	(0.024)	(0.027)	(0.029)
Area: Imm. share	0.115	0.146	0.214	0.254
	(0)	(0.001)	(0.001)	(0.001)
Labour income	110800	96220	92700	95620
	(224.1)	(205.6)	(187.2)	(169.5)
Area: Share high ed.	0.254	0.219	0.16	0.135
	(0.001)	(0.001)	(0.001)	(0)
Municipal sector	0.639	0.545	0.426	0.201
	(0.002)	(0.002)	(0.002)	(0.002)
Ind.: Manufacturing	0.099	0.146	0.241	0.473
	(0.001)	(0.002)	(0.002)	(0.002)
Immigrant	0.037	0.188	0.333	0.486
	(0.001)	(0.003)	(0.004)	(0.004)
Area: sick leave	16.03	17.91	22.24	24.22
	(0.027)	(0.033)	(0.038)	(0.04)
N workers at plant	1505	1860	2273	2886
	(13.03)	(16.94)	(18.98)	(19.86)
Age	32.57	35.24	36.51	36.55
	(0.045)	(0.053)	(0.055)	(0.053)
Area: earnings	134600	130900	122600	119000
	(91.19)	(92.72)	(75.25)	(65.08)
Earnings rank, plant	0.599	0.55	0.544	0.5
	(0.001)	(0.001)	(0.001)	(0.001)
Edu. field: Health	0.311	0.165	0.082	0.021
	(0.002)	(0.002)	(0.001)	(0.001)
Married	0.499	0.497	0.488	0.437
	(0.002)	(0.002)	(0.002)	(0.002)

	Quartile 1	Quartile 2	Quartile 3	Quartile 4
Commuter	0.137	0.104	0.088	0.098
	(0.002)	(0.001)	(0.001)	(0.001)
Edu. field: Sci/eng	0.132	0.169	0.209	0.324
	(0.001)	(0.002)	(0.002)	(0.002)
Tenure	1.368	1.445	1.477	1.563
	(0.006)	(0.006)	(0.006)	(0.006)
Days in hospital	0.266	0.651	0.69	0.933
	(0.009)	(0.02)	(0.022)	(0.035)
Ind.: Public admin	0.028	0.026	0.022	0.031
	(0.001)	(0.001)	(0.001)	(0.001)
Mun. pop. density	710.8	757.3	827.1	843.1
	(1.856)	(1.734)	(1.475)	(1.401)
Divorced	0.063	0.105	0.134	0.133
	(0.001)	(0.001)	(0.002)	(0.002)
Ind.: Construction	0.016	0.036	0.055	0.105
	(0.001)	(0.001)	(0.001)	(0.001)
Children 7-18	0.41	0.428	0.397	0.353
	(0.003)	(0.003)	(0.003)	(0.003)
Edu. field: Admin/law	0.229	0.183	0.115	0.064
	(0.002)	(0.002)	(0.001)	(0.001)
Children<7	0.26	0.236	0.22	0.18
	(0.003)	(0.002)	(0.002)	(0.002)
Gothenburg	0.734	0.783	0.857	0.874
	(0.002)	(0.002)	(0.002)	(0.001)
Ind.: Business	0.104	0.086	0.077	0.053
	(0.001)	(0.001)	(0.001)	(0.001)
Ind.: Utilities	0.045	0.055	0.078	0.1
	(0.001)	(0.001)	(0.001)	(0.001)
Edu. field: Services	0.047	0.058	0.049	0.038
	(0.001)	(0.001)	(0.001)	(0.001)
Edu. field: Teacher	0.09	0.049	0.013	0.005
	(0.001)	(0.001)	(0.001)	(0)
Ind.: Education	0.027	0.033	0.029	0.011
	(0.001)	(0.001)	(0.001)	(0)
Ind.: Primary	0.003	0.005	0.005	0.008
	(0)	(0)	(0)	(0)

Note: Quartiles ranked according to causal forest estimated treatment effects, with Q1 containing those estimated to be least affected and Q4 those estimated to be most affected. Colours are assigned according to how strongly the average value of the variable in the quartile deviates from its grand mean.

Essay II. Consequences of Job Loss for Routine Workers

I am grateful to IFAU for data access and to Stefan Eriksson, Georg Graetz, Adrian Adermon, Michael Böhm, Marcus Eliason, Duncan Roth, conference participants at AIEL 2022 and seminar participants at Uppsala University for their helpful suggestions and comments. Financial support from Vetenskapsrådet, grant number 2018-04581, is gratefully acknowledged.

1. Introduction

Over the last several decades, automation of tasks that had previously required human labour has taken place at a rapid pace. While this has led to increased labour productivity (Graetz and Michaels, 2018), concerns as to the effects of labour-replacing technology on worker welfare and income inequality have been raised both within academia and in the broader public debate (Acemoglu and Autor, 2011). The consensus in the literature is that automation has indeed been a contributing factor to increased income inequality in developed countries in recent decades, with the main effect operating through its impact on the occupational distribution of the workforce. Machines have tended to replace workers in middle-skill, middle-wage manufacturing and clerical occupations, while complementing labour in high-skilled managerial and professional positions. The share of employment in low-wage service jobs, which have been relatively unaffected by automation, has also increased. Overall, this has led to the workforce being increasingly polarised in terms of occupational skills and wages (Autor, et al., 2003; Goos et al., 2016). The reduction in routine employment has likely involved a large number of involuntary job separations, as firms have laid off workers whose input is no longer required in production. However, evidence on how routine workers fare following involuntary job loss has been scarce. Indeed, most previous work has focused on the aggregate labour market effects of technological change rather than the impact on exposed workers. Nevertheless, there are studies suggesting that workers in declining occupations have suffered from reduced employment and earnings (Edin et al., 2019) and that workers in routine occupations have seen lower wage growth than other worker categories (Cortes, 2016). Theory suggests that displaced routine workers should do worse than displaced non-routine workers, as they are likely to have a harder time finding a new job that fits their occupation-specific skills, in addition to facing the loss of good employer-employee matches, firm-specific human capital and rents as all workers do. Indeed, the direct exposure of involuntarily displaced routine workers suggests they could be among the biggest losers of automation and technological change. The magnitude of their losses might provide an approximate upper bound on the detrimental effects of labour-replacing technology. In this paper, I seek to establish whether routine workers are more affected by a common type of involuntary job displacement, namely establishment-level closures and mass layoffs.

The evidence on how establishment shutdowns and mass layoffs affect worker outcomes is extensive and overwhelmingly negative. Since the pioneering paper by Jacobson et al. (1993), studies have almost invariably found that job loss has severe impacts on workers' subsequent employment, earnings and even health (Sullivan and von Wachter, 2009; Davis and von Wachter, 2011). Worker outcomes do not regain the levels of comparable controls who avoid losing their jobs for many years, resulting in a seemingly permanent scarring effect. These results hold in practically all countries where this question has been investigated; Eliason and Storrie (2006) show that Swedish workers do not recover in terms of labour market outcomes even 12 years after losing their jobs. There is evidence that displaced workers fare worse when demand for either labour in general or for their particular occupation- or industry-related skills is low due to aggregate economic conditions (Davis and von Wachter, 2011), local occupation-specific labour demand (Galaasen and Kostol, 2018) and import competition (Dauth et al., 2021). This suggests that if some occupations experience rising demand due to complementary technological change, while others decline due to automation, the experiences of workers in these occupations should be different following job loss. Furthermore, there are indications that displaced workers whose skills are not in demand suffer larger losses than their peers (Nedelkoska et al., 2022). The first comparison of the post-layoff outcomes of routine and non-routine workers is conducted in a recent paper by Blien et al. (2021), whose results point to substantial penalties for routine workers in terms of earnings and employment, but only insignificant effects on their wages.

In order to assess differences in the size of routine and non-routine workers' post-layoff losses, I use a standard difference-in-difference event study approach. Displaced individuals' labour market outcomes at different time points preceding and following layoff are compared to those of similar non-displaced workers. Routine and non-routine workers who lose their jobs are compared to corresponding groups of non-displaced peers. Matching on a large set of characteristics, including age, gender, education, tenure, size of closing establishment, size of local labour market as well as broad industry and occupation categories ensures that the groups of displaced and non-displaced workers are observationally comparable. Detailed Swedish matched employer-employee data enable me to identify all those who lost their jobs in plant shutdowns or mass displacement events during the 1997-2014 period, although occupational information is missing for a fraction of individuals. Individual worker outcomes are tracked for ten years following the year an establishment shuts down or experiences a mass layoff. The aim is to be as representative of all displaced workers as possible, including small workplaces (5-49 employees), older workers aged 51-62 and all public sector workers (including civil servants).

The results show that layoff penalties are significantly more adverse for routine workers than for non-routine workers. Their labour income falls by 20 percentage points more than that of non-routine laid off workers in the year following displacement, and remains significantly lower for eight years. This drop is mostly due to lower re-employment probabilities for displaced routine workers; the probability of not being employed in the year following displacement is 11 percentage points higher for routine workers than for their non-routine counterparts. The monthly wages of laid off routine workers also drop five log points more than what is the case for comparable non-routine workers.

Seen from another perspective, routine workers spend 90 additional days in unemployment in the first post-displacement year. Overall, the evidence suggests that workers exposed to automation suffer greatly when they are displaced from their jobs. The estimated effects are larger than those found in studies that have considered individuals in routine or otherwise declining occupations in general, without focusing specifically on mass layoffs (Cortes, 2016; Edin et al., 2019). A share of these losses may be due to losses of occupation and industry-specific human capital, as routine workers are more likely to find new employment outside of their original occupation and sector. This view is reinforced by the fact that displaced routine workers are likely to move to lower-paying industries and to end up with lower earnings compared to other workers in their new occupation. Switchers from routine to non-routine occupations do worse in terms of earnings than those who continue doing routine work. This is in line with earlier results on costs of occupational mobility increasing in task and skill distance (Cortes and Gallipoli, 2018; Robinson, 2018), but contrasts with Cortes' (2016) findings that switchers from routine to non-routine cognitive occupations see wage increases. This difference could be due to a higher prevalence of involuntary switchers among displaced workers.

The analysis is important for establishing the external validity of the findings of Blien et al. (2021), as it is conducted using high-quality data from another country. Furthermore, the number of days spent in unemployment is studied as an outcome, providing more concrete evidence as to whether reductions in employment and earnings are involuntary. Finally, unlike Blien et al. (2021), I find significant negative effects on displaced routine workers' wages and show that they are more likely to transition across industries.

The remainder of this paper is structured as follows. Section 2 describes the data used, explains how routineness is defined, provides descriptive statistics for displaced and non-displaced workers and covers the labour market outcomes included in the study. The empirical model estimated is presented in Section 3 and the results, along with robustness checks, heterogeneity analysis and a discussion of mechanisms are shown in Section 4. Section 5 concludes.

2. Data

2.1 Selection of Displaced and Control Samples

I use a rich micro-level dataset created by Statistics Sweden which contains information on all Swedish employment relationships. Data on occupation are collected in the Wage Structure Statistics dataset and are available for all public sector and a large sample (about half) of private sector workers. The probability of a private firm being sampled is determined by its size, with large firms overrepresented. If there is no information on a worker's occupation in the current year, it is imputed using reported occupations in the three preceding years on the condition that the worker has remained at the same establishment. Years before 1996 are excluded because converting the old occupational codes to the new system is very difficult. In all cases, the focus is on a worker's main place of employment during a given year. This is defined as the establishment where the worker had his or her highest source of earnings that year.

The displacement and control groups of establishments are defined based on the change in the number of workers for whom they are the main place of employment. Shutdowns are defined as cases where an establishment ceases to be the main workplace of any worker. The establishment is required to exist in year t_{-1}, to be the main workplace of at least one individual in the event year t_0, and to no longer be the main workplace for anyone in t_1, the year after the event. If the number of individuals who have their main place of employment at an establishment falls by at least 80 percent from t_{-1} to t_1, this is classified as a mass layoff event.[1] Events where more than 30 percent of the displaced workers end up at other establishments in the old workplace's firm or in the same unique establishment at a different firm in t_1 are excluded from both the displacement and control groups. This is standard in the literature because these cases are likely to be firm mergers, acquisitions or reorganisations rather than real job displacement events (Hethey-Maier and Schmieder, 2013). Establishments with fewer than five workers in t_{-1} are also excluded in order to reduce the possibility of individual worker characteristics having a large impact on overall plant performance. This size restriction is among the most permissive used in the literature. The control group of establishments consists of those that had at least five employees in t_{-1} and did not experience a shutdown or mass layoff from t_{-1} to t_0 or from t_0 to t_1.

Workers who had their main place of employment at a closing establishment in the year t_{-1} immediately preceding the shutdown or layoff event are categorised as displaced. Early leavers are thus captured in the displaced sample as there is no requirement that individuals work at the shutting establishment in t_0, the year of shutdown or layoff. Also, workers who stay at their old workplaces following a mass layoff are included in the displaced sample to avoid the issue of selectiveness in terms of who gets laid off. To ensure that the individuals studied have a sufficiently strong connection to the shutting establishment, they are required to have a tenure of at least two years (defined as having their main place of work at the shutting establishment in the years t_{-2} and t_{-1}). This restriction is less stringent than what is typically used in the literature and aims to minimise the number of workers with a strong degree of

[1] Similar cutoffs are used by e.g. Davis and von Wachter (2011). A number of studies also include events such as employment decreases of 30 percent or more. However, it is problematic to use such a cutoff when including small establishments, as changes in the establishment not related to mass layoffs might lead to such employment shifts.

attachment to the closing establishment who are excluded. The control pool consists of those who were employed at a control establishment in the year t_{-1} and had tenure of at least two years. There are no conditions on what happens to control workers or their establishments beyond t_{-1}; it is possible for them to themselves become displaced at a later point in time. This avoids the downward bias on displacement loss estimates that appears when the control group is defined conditional on never being displaced (Krolikowski, 2018). Workers who are younger than 22 or older than 62 in the year prior to layoff are dropped from both the displaced and control groups. The sample thus includes older workers, who are sometimes excluded in other studies. Older workers are not followed after they reach the age of 65, as this is the typical retirement age. In order to ensure that the workers considered are at least somewhat consistently attached to the labour market, the sample is limited to those who earn at least three times the tenth percentile-level blue-collar monthly wage in each of the years t_{-4} through t_{-1}.[2] This restriction also entails dropping workers who are not continuously observed in the Swedish registry data in the four years prior to the real or placebo displacement event. As an additional safeguard against including individuals only tenuously attached to the labour market, workers who were registered as unemployed for 183 days or more in any of the years t_{-4} through t_{-2}, or for 330 days or more in the year t_{-1} are excluded. The more liberal restriction on the year t_{-1} aims to exclude as few early leavers as possible; this concern arises because days spent in unemployment begin rising for displaced workers already in t_{-1}. Workers who are not observed in both the years t_0 and t_1 are also removed from the sample because their post-layoff outcomes are not known. Finally, individuals for whom occupational data are missing even after imputation are excluded as the routineness of their jobs cannot be determined. This final condition is the most restrictive, as occupational data are missing for 52 percent of eligible displaced and 28 percent of eligible controls. After restrictions are imposed, the eligible sample of displaced workers consists of 84,896 individuals who lose their jobs in 4,866 shutdown or mass layoff events.

2.2 Routineness Definition

Routineness is defined based on the Dictionary of Occupational Titles (DOT), as has been standard in the routine-biased technological change literature since the seminal study by Autor et al. (2003). The US occupations whose task intensities are determined using the DOT are translated to the ISCO-88 international classification, which is in turn matched to corresponding Swedish occupations. Routineness is measured as the sum of an occupation's intensities

[2] The tenth percentile of blue-collar monthly wages is used as a measure of the "minimum wage", as minimum wage legislation is absent in Sweden. Wage levels are instead agreed through collective bargaining between unions and employer organisations.

in tasks that are routine cognitive ("set[ting] limits, tolerances, or standards" according to the DOT) and routine manual ("finger dexterity" according to the DOT). The sum of intensities in these two task categories is normalised by the occupation's total intensity in all tasks. This provides a measure of the share of routine tasks in the total number of tasks involved in the occupation, which should be a good measure of its exposure to automation. In the main specification, the cutoff for being classified as routine is set at the upper quartile of routineness among workers displaced in 2005, in the middle of the studied period. One quarter of the workers displaced in 2005 are therefore classified as routine and three quarters as non-routine. This results in 15 three-digit occupations being categorised as routine and 81 as non-routine. This grouping of routine and non-routine occupations is also applied for those displaced in other years. As expected, routine occupations consist exclusively of machine operating, clerical, elementary, and some crafts jobs, while non-routine occupations are typically managerial, high-skilled or service jobs. Setting a high cutoff for routineness makes it more likely that the occupations classified as such are indeed exposed to labour-replacing technological change; nevertheless any threshold is somewhat arbitrary and alternative definitions are tested. These are setting the cutoff for routineness at the occupation of the displaced worker with median routineness in 2005 and dropping the occupations of the middle two quartiles of workers entirely to only include high-routine and low-routine occupations. These alternative definitions do not qualitatively affect the results.

2.3 Descriptive Statistics and Matching

Descriptive statistics for displaced and non-displaced workers in terms of routineness, demographics, education, occupation, industry and location are shown in Table 1. All individual characteristics are defined based on the year t_{-1} preceding the year t_0 in which the mass layoff takes place. This should reduce the risk of changes immediately related to plant closure having an effect. The first two columns of Table 1 show that displaced workers tend to be in slightly more routine occupations than non-displaced ones. This is mainly explained by an extreme overrepresentation of manufacturing workers among the displaced; almost half of them were in a manufacturing job prior to layoff. This affects the occupational composition of the displaced, which is skewed towards operators, assemblers and crafts workers. On the other hand, it is very rare for displaced workers to be found in typically public sector industries, such as education, health and public administration. High-skilled professional workers are also underrepresented.

Because of these discrepancies, I use propensity score matching to make the two groups of workers more comparable. In the main specification, displaced and control individuals are matched on routineness, age, gender, education level, tenure, broad industry, establishment size, a measure of their municipality's urban character, and earnings in periods t_{-4} through t_{-1}. Because the analysis focuses on occupations, workers are matched within broad one-digit occupational groups. To avoid comparing trajectories of workers who were displaced in different years, matching is done within cohorts defined by the calendar year of the real or placebo event. Each displaced worker is assigned one match from the pool of controls with replacement. Workers whose propensity scores lie outside of the common support region where the propensity score distributions of the displaced and controls overlap are trimmed away. As can be seen in the top panel of Appendix Figure A1, the propensity scores of the unmatched controls skew heavily towards zero, while those of the displaced are more spread out. However, as the size of the control pool is much larger than the number of displaced, over 99 percent of the displaced workers are within the common support region. Good matches are available for practically all displaced workers, as can be seen in the bottom panel of Appendix Figure A1. The propensity score distributions among matched displaced and controls overlap almost perfectly.

Descriptive statistics for the matched samples are presented in the last two columns of Table 1. Matching within broad occupational groups ensures perfect balance in that dimension. The matched sample of controls is also very similar to the displaced in terms of routineness, demographics, education level, industry, and municipality type. To ensure robustness of results, I test alternative matching specifications where either the entire sample is used without any restrictions, or matching is done based on the covariates listed above, but excluding pre-period earnings. Neither of these other specifications produces results qualitatively different from those given by the main specification. The pre-displacement values of the outcome variables on which I do not match (employment probability, monthly wages and days of unemployment) are presented in Table A1 in the Appendix. Table A2 contains the post-matching characteristics for routine and non-routine workers (as defined by the cutoff used in the main specification) separately.

Table 1. *Descriptive statistics for the matched and unmatched samples of controls and displaced*

	Controls (Unmatched)	**Displaced** (Unmatched)	**Controls** (Matched)	**Displaced** (Matched)
N individuals	1,035,499	84,896	65,069	84,325
Routine intensity	0.51	0.58	0.58	0.58
Year t_{-1}	2005.0	2004.3	2004.3	2004.3
Age	45.3	43.3	43.3	43.3
Tenure	6.2	5.6	5.5	5.6
Female	0.54	0.37	0.37	0.37
Immigrant	0.10	0.11	0.11	0.11

Education level (percentages)				
Less than compulsory	4.14	7.00	7.16	6.99
Compulsory, 9 years	7.97	12.91	13.04	12.91
High school, 2 years	31.15	32.68	33.07	32.70
High school, 3 years	16.71	21.74	21.16	21.74
Some post-secondary	13.87	12.73	12.22	12.74
University	24.29	12.14	12.52	12.14
PhD	1.86	0.79	0.82	0.78
Occupations (percentages)				
Officials & Managers	5.55	6.88	6.87	6.87
Professionals	24.16	12.47	12.47	12.47
Technicians	17.98	17.35	17.38	17.38
Clerks	9.36	11.45	11.46	11.46
Service & Sales	20.60	10.14	10.12	10.12
Crafts	6.60	11.34	11.34	11.34
Operators & Assemblers	10.79	23.70	23.80	23.80
Elementary Occupations	4.96	6.67	6.57	6.57
Industries (percentages)				
Primary	0.74	0.40	0.46	0.40
Manufacturing	21.98	48.77	48.17	48.88
Construction	2.52	2.21	2.22	2.21
Utilities & telecom	6.42	9.71	9.88	9.68
Wholesale & retail	6.96	10.36	10.67	10.35
Business services	10.63	17.09	16.92	17.03
Health, social work	29.16	7.53	7.59	7.53
Education	14.64	1.52	1.61	1.51
Public administration	6.96	2.42	2.47	2.42
Type of municipality (percentages)				
Rural municipalities	14.47	15.64	15.88	15.66
Commuter municipalities	4.45	5.62	5.64	5.51
Towns	16.28	15.00	14.88	15.05
Other cities	33.38	32.61	31.87	32.66
Suburbs of 3 largest cities	10.52	10.04	10.43	10.06
3 largest cities	20.91	21.09	21.29	21.07

Pre-period earnings (SEK thousands)				
t_{-1}	323	334	333	334
t_{-2}	316	321	323	321
t_{-3}	306	308	311	309
t_{-4}	293	294	296	295

Note: Characteristics evaluated in year t_{-1} unless stated otherwise. Unmatched control group consists of 5% random sample of the eligible control pool. One-to-one propensity score matching with replacement implemented based on characteristics listed in the table. Propensity scores estimated using logit. Matched control sample statistics weighted by the number of times a control worker was drawn as the best match for a displaced worker. Sum of matched control weights is 84,325.

2.4 Outcomes Studied

The annual earnings outcome is normalised by the mean of the worker's earnings in t_{-4} through t_{-1}. This provides an individual baseline for each worker and makes the size of the estimates independent of absolute differences in pre-period earnings and wages of routine and non-routine workers. Annual earnings are measured before income tax. The employment outcome is a dummy for earning at least three times the tenth percentile-level blue-collar monthly wage within a given year.[3] Unemployment is measured as the number of days the individual is registered as unemployed or taking part in a labour market programme at the Public Employment Service. In Sweden, one must register as unemployed in order to receive benefits, meaning that instances of unemployed individuals abstaining from registering should be minimised. The days of unemployment measure represents the total number of days, including weekends and holidays, rather than only working days. Wages are the worker's monthly wages at their main workplace, measured in the second half of the year. Data on wages are available only for workers who were sampled into the Wage Structure Statistics that year. This is the same sample as the one from which occupational information is obtained (it contains all public sector workers and about half of private sector workers, with large firms overrepresented). This means that wage data are missing for many individuals in at least some years, while the other outcomes are always observed for the population of displaced and control workers.

3. Empirical Specification

An event study approach typical for the literature is used. The main effect of job displacement is estimated first using a differences-in-differences model:

[3] This is analogous to the definition of employment for the purposes of determining attachment to the labour market in the pre-period. The tenth percentile of blue-collar monthly wages is once again used as a measure of the "minimum wage", as minimum wage legislation is absent in Sweden.

$$y_{it} = \sum_{\tau=-4,\tau\neq-1}^{10} [\alpha_\tau I(t = t_0 + \tau) + \beta_\tau I(t = t_0 + \tau) \times D_i] + \lambda_i + \mu_t + \varepsilon_{it} \quad (1)$$

The model given by equation (1) regresses a labour market outcome, such as earnings or wages, on a set of dummies $I(t = t_0 + \tau)$ for years relative to the year of real or placebo displacement, which is indexed by t_0. The coefficients on year t_{-1} have been normalised to zero. The main effects of interest are given by the set of β_τ which measure the size of the interaction effect between year dummies and actual displacement. Individual fixed effects λ_i are included to remove influences from time-invariant individual characteristics, which can affect the estimates as the panel of workers is unbalanced. General economic conditions in a given year are controlled for by calendar year dummies μ_t.

The main specification is based on Equation (1), but adds a full set of routine-time-to-event and routine-displacement indicators. A full set of routine-calendar year interactions is also included to control for general trends in routine labour market outcomes in the economy, which is necessary as the panel is not fully balanced (results using a fully balanced sample of individuals who are observed during the entire t_{-4} to t_{10} period are presented in the Appendix for comparison). The following equation results:

$$\begin{aligned} y_{it} = \sum_{\tau=-4,\tau\neq-1}^{10} & [\alpha_\tau I(t = t_0 + \tau) + \beta_\tau I(t = t_0 + \tau) \times D_i + \delta_\tau I(t = t_0 + \tau) \\ & \times R_i + \gamma_\tau I(t = t_0 + \tau) \times D_i \times R_i] + \lambda_i + \mu_t + \mu_t \times R_i + \varepsilon_{it} \end{aligned} \quad (2)$$

Now, the effects of displacement on non-routine workers relative to non-routine workers who are not displaced are given by the set of β_τ. The effects of displacement on routine workers relative to non-displaced routine workers are given by $\beta_\tau + \gamma_\tau$, with γ_τ capturing any differences in the displacement penalty between routine and non-routine workers. In the figures below, the non-routine series plot the estimates β_τ, while the routine series show $\beta_\tau + \gamma_\tau$. Standard errors are clustered at the level of the t_{-1} establishment in all cases.

4. Results

4.1 Post-Layoff Outcomes of Routine and Non-Routine Workers

The baseline estimated effects of job loss on real earnings from Equation (1) are shown in Figure 1. Earnings evolve in a very similar fashion for displaced and non-displaced workers through t_{-1}. The relative earnings of the displaced then decrease somewhat in t_0 (the last year the closing establishment is observed) before falling sharply to 25 percent less than the pre-displacement earnings level in t_1. While there is some recovery in the following years, displaced workers' earnings never regain the trajectories of their non-displaced peers, remaining 13 percent lower in t_{10}. This pattern is similar to what previous studies have found (Jacobson et al., 1993; Eliason and Storrie, 2006; Davis and von Wachter, 2011). This is in spite of some differences regarding sampling restrictions, suggesting that they do not have a qualitative effect on the findings.

The results of the main earnings specification as estimated by Equation (2) are presented in Figure 2 (the point estimates are also shown in Table A3 in the Appendix). Although the trajectories of earnings for routine and non-routine workers follow each other closely in the period up to displacement, they diverge clearly in t_0. While non-routine workers lose 20 percent of their pre-displacement earnings in t_1, their worst post-displacement year, for routine workers the corresponding share is 39 percent. However, the earnings of laid off routine workers converge with those of their non-displaced counterparts more quickly than those of non-routine workers. This means that the gap between the two groups of displaced workers narrows over time. Nevertheless, the additional penalty suffered by routine workers remains statistically significant for eight years after establishment closure. Cumulatively, non-routine workers are estimated to lose 1.26 times the amount of a year's worth of pre-displacement earnings over the t_0 to t_8 period. Routine workers are estimated to lose 2.22 times worth of their pre-displacement annual labour income over the same time frame. The convergence seems to be driven by the fact that many non-displaced routine workers have disadvantageous earnings trajectories; their real earnings are only 5.7 percent higher in t_{10} than during the t_{-4} to t_{-1} period, while the real earnings of non-routine control workers grow by 20 percent over this timeframe.[4] Indeed, by t_{10} displaced non-routine workers are estimated to have higher earnings relative to the baseline period than non-displaced routine workers.

[4] See Table A1.

Figure 1. *Baseline estimate of effects of job loss on earnings*

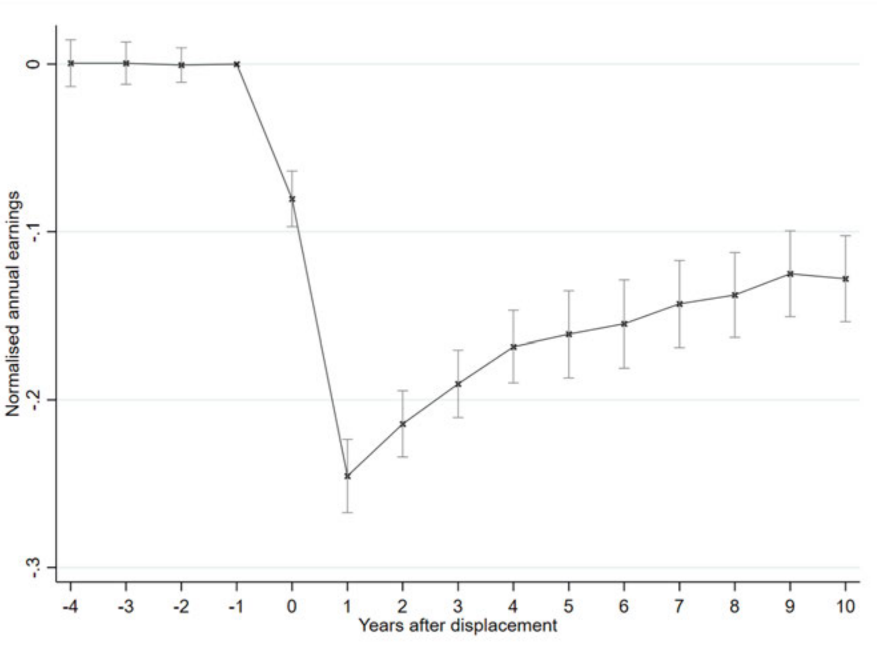

Note: The baseline period is the year before displacement, t_{-1}. Outcome of displaced relative to matched control group in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

Figure 2. *Estimated effects of job loss on earnings for routine and non-routine workers, relative to non-displaced workers in the respective category*

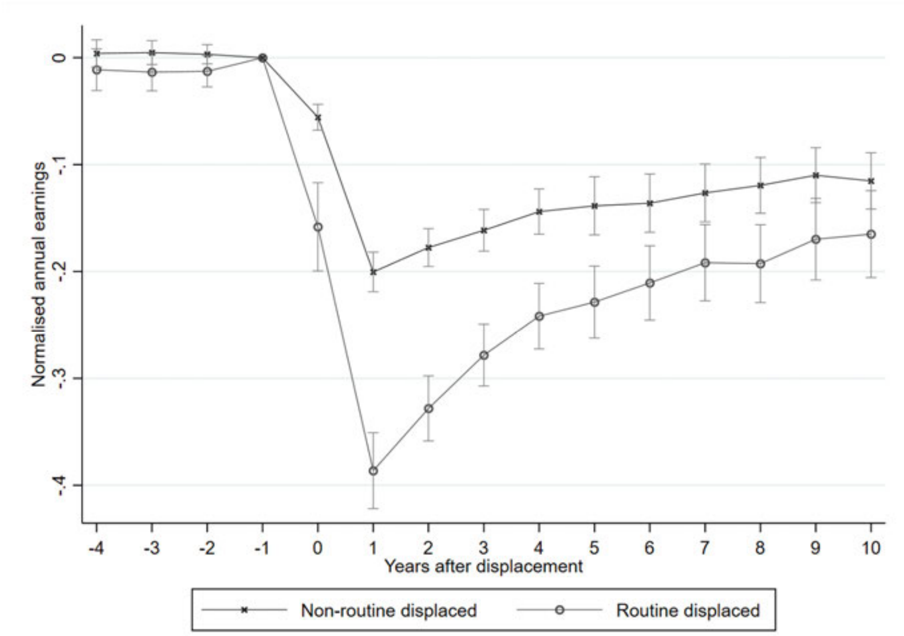

Note: The baseline period is the year before displacement, t_{-1}. Outcome of routine and non-routine displaced relative to routine and non-routine matched controls in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

A breakdown of annual earnings losses reveals substantial adverse effects of displacement on both the probability of being employed and on wages conditional on employment. Figure 3 shows the estimated effect of displacement on the employment probability for routine and non-routine workers (the baseline employment effect for all workers is plotted in Figure A2 in the Appendix). All workers are employed by construction in the four years from t_{-4} through t_{-1} and there is thus no difference between routine and non-routine workers in this regard. However, by t_1 displaced routine workers are 11 percentage points less likely to be employed than displaced non-routine workers. This difference is persistent, and even though it narrows over time, is statistically significant through the fifth post-layoff year. Just like in the case of earnings, neither group of workers fully recovers from the shock of losing their jobs. In the case of monthly wages, results for which are shown in Figure 4, estimates are somewhat noisy because wage data are not available for the full sample of workers each year (the baseline wage effect can be seen in Figure A2). They do however indicate that routine workers suffer much more following displacement, suffering a 6.9 log point drop in wages in t_1, while non-routine workers only see wages drop by 1.7 log points. The difference remains significant for the first four post-layoff years. It seems that routine workers' wages converge more quickly to the level of their non-displaced peers than is the case for non-routine workers, whose wages do not seem to converge at all. However, it is difficult to draw any definitive conclusions about this as the point estimates for different years are noisy and not statistically distinguishable in most cases.

Figure 3. *Effects of job loss on the probability of being employed for routine and non-routine workers, relative to non-displaced workers in the respective category*

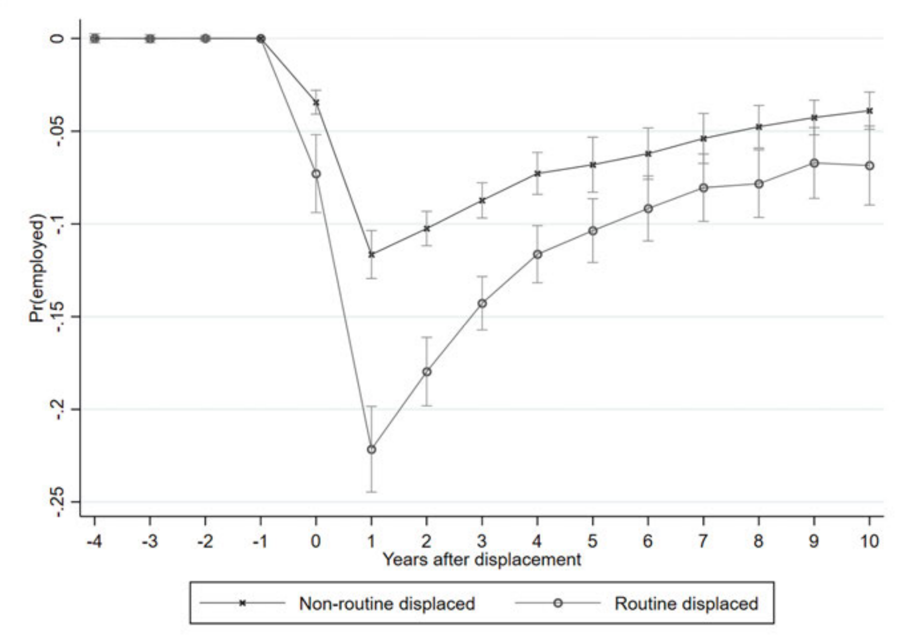

Note: Baseline period is the year before displacement, t_{-1}. Outcome of routine and non-routine displaced relative to routine and non-routine matched controls. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

Figure 4. *Effects of job loss on log monthly wages (conditional on being employed) for routine and non-routine workers, relative to non-displaced workers in the respective category*

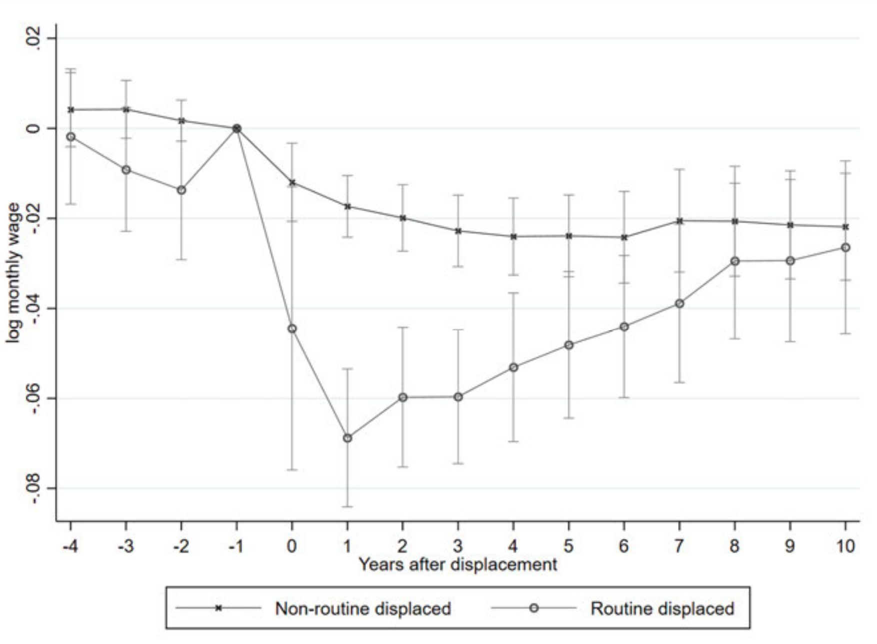

Note: Baseline period is the year before displacement, t_{-1}. Outcome of routine and non-routine displaced relative to routine and non-routine matched controls. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

Finally, I turn toward an alternative way of measuring adverse labour market outcomes, namely the number of days in a year registered as unemployed. According to this metric, routine workers also suffer more following displacement than non-routine ones do, as can be seen in Figure 5 (Figure A2 shows the average unemployment effects of displacement). The largest unemployment effects are observed in the year t_1, when non-routine displaced workers spend 39 more days in unemployment than their non-displaced counterparts. At the same time, displaced routine workers experience 88 additional days of unemployment. The difference in time spent unemployed is persistent and remains statistically significant, although quantitatively smaller, until the sixth post-displacement year. By this time, displaced routine workers have on average spent a total of 307 additional days in unemployment, compared to 126 days for displaced non-routine workers.

Figure 5. *Estimated effects of job displacement on days spent in unemployment for routine and non-routine workers, relative to non-displaced workers in the respective category*

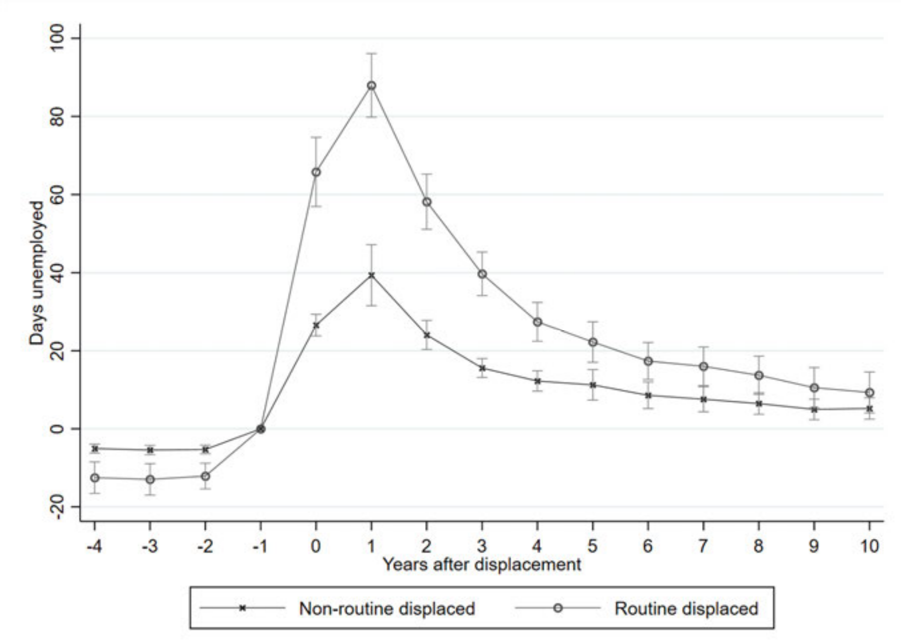

Note: The baseline period is the year before displacement, t_{-1}. Outcome of routine and non-routine displaced relative to routine and non-routine matched controls in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

4.2 Robustness Checks

To ensure that the results are not sensitive to the definition of routineness used or the type of matching employed, several robustness checks are employed. While the definition of routineness used in Section 4.1 makes it more likely that truly routine occupations are classified as routine due to the stringent cut-off used, this definition might be too narrow. For this reason, Equation 2 has been re-estimated using an alternative routineness definition where the workers are split by median routineness instead of classifying only the most routine quartile as routine.[5] According to this definition, 43 three-digit occupations are routine and 53 are non-routine. Another possibility is that the presence of occupations that are close to one another in terms of routineness on both sides of the threshold attenuates the results. To make sure that this is not the case, Equation 2 is estimated using only those individuals whose occupations were

[5] Just like in the main specification, displaced workers who are in their t_{-1} period in 2005 are ordered by routineness, but the split is at the occupation of the median worker rather than the occupation of the third quartile worker.

either in the top or bottom quartiles of displaced workers ordered by routineness in t_{-1}. This leaves the 15 occupations classified as routine in the main analysis and 30 low-routineness occupations.

These alternative definitions yield results almost identical to those given by the main specification, as can be seen in Figure 6. The top panel plots the baseline estimates of the effects of displacement on earnings, reproducing Figure 2. The panel on the bottom left shows the results when the median is used as the threshold for the routine category and the panel on the bottom right shows the results when only the top and bottom quartiles of routineness are included. Using a less stringent definition of routineness reduces the size of the estimated routine penalty in the years immediately following layoff. Also, limiting the non-routine sample to those in the lowest quartile of routineness leads to slightly larger routineness penalty estimates. These apparent differences are in line with what theory predicts. The routine penalty estimates and the post-layoff earnings trajectories are very similar in the three specifications, confirming that the way routineness is defined is not of key importance for the results. A final alternative specification where post-layoff outcomes are plotted for each of the four routineness quartiles separately is shown in Figure A3 in the Appendix. The fourth routineness quartile, which contains the workers classified as routine in the main specification, does clearly worse than the other three quartiles. The differences between the first, second and third quartiles are not as clear. In t_1, workers from the second and third routineness quartiles appear to suffer larger penalties than those in the first quartile, but the trajectories of these groups converge over the medium and long run.

Graphs corresponding to Figure 6 for the employment and monthly wage outcomes are presented in Figures A4 and A5 in the Appendix. In the case of employment, the differences between the different definitions are small, although there are indications that the routine penalty is smaller if all workers with above-median routineness are categorised as routine. However, with this definition of routineness, the routine wage penalty becomes insignificant in all years except for t_1. However, the confidence intervals are wide enough to contain the estimates from the main specification. Results for days of unemployment using the different routineness definitions are shown in Appendix Figure A6. The definitions give similar results, except for a somewhat smaller routine penalty in t_1 when the median cutoff is used.

In addition to estimating Equation 2 using a sample matched on covariates and earnings in t_{-4} through t_{-1}, I estimate it in turn using the full unmatched sample and a sample matched only on covariates, but not pre-period earnings. These alternative samples of workers give results very similar to those obtained using the preferred sample. Their results are presented in Figure A7 in the Appendix. Point estimates of earnings penalties when the unmatched sample is used are also provided in Appendix Table A3 for all three definitions of routineness.

Finally, Equation 2 has been re-estimated using a fully balanced panel, leaving only those workers who are observed in the Swedish registry data and are younger than 65 years of age in each of the years t_{-4} through t_{10}. This entails reducing the sample to shutdowns and mass layoffs that took place in 1997-2006, as data for years after 2016 is not available. Also, workers older than 52 years of age in t_{-1} are excluded. The results of this exercise for the outcomes of earnings and unemployment are shown in Appendix Figure A8. Using the fully balanced panel reduces the size of penalties immediately following layoff for both routine and non-routine workers (the differences from the full panel estimates are rarely statistically significant), but has no effect on penalty estimates for later years. The routine penalty remains large and statistically significant.

Figure 6. *Estimates of displacement earnings penalties using different definitions of routine and non-routine occupations*

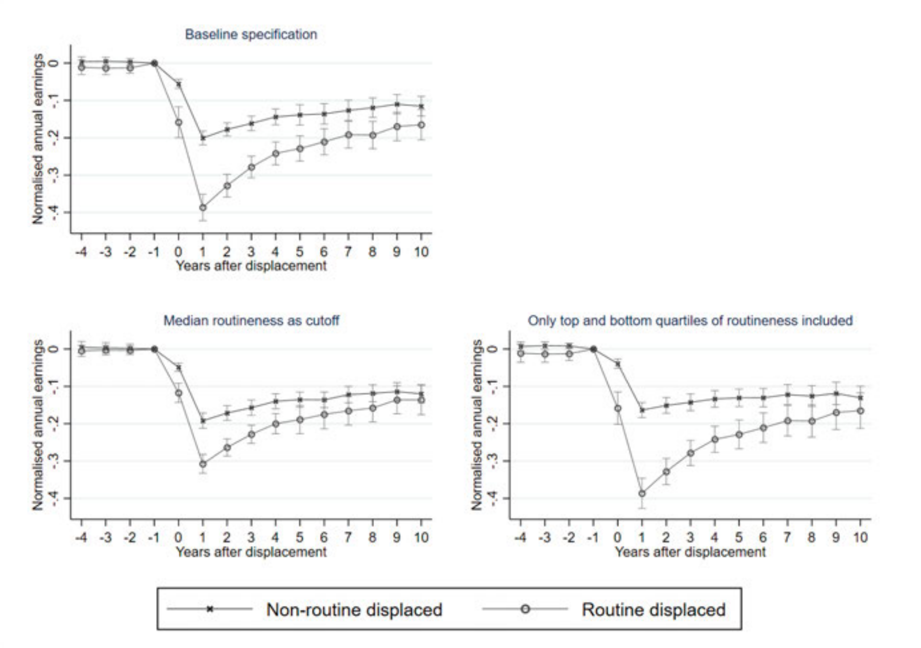

Note: The baseline period is the year before displacement, t_{-1}. Outcome of routine and non-routine displaced (according to different definitions of routineness) relative to routine and non-routine matched controls in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

4.3 Heterogeneity in Routine Penalties

I consider heterogeneity in the size of routine penalties for workers with different levels of education, by sector, and within broad occupational categories. The top panel of Figure 7 shows post-displacement earnings trajectories among workers with high school education or less and among workers with more than a high school education.[6] Among less educated workers, the trajectories of routine and non-routine workers are very similar to those among the full displaced sample. Among highly educated workers, initial losses among the non-routine group are initially somewhat smaller than in the full sample, but routine workers' losses are not. The size of the penalty for routine highly educated workers decreases before seeming to actually increase again at the very end of the period studied, but this is likely to be an artefact of the small sample size, as confidence intervals are very wide.

The middle panel of Figure 7 shows earnings trajectories of workers who are displaced in the manufacturing[7] and services sectors. The results for manufacturing are similar to the findings for the full sample, albeit with indications that non-routine workers who are displaced in manufacturing do slightly worse. In the service sector, penalties for both routine and non-routine workers are lower. Also, it seems that convergence of routine workers' losses to the level experienced by non-routine workers is quicker. The difference between the two groups only remains significant through t_2 and the point estimates for t_6 and t_7 are almost identical for the groups of routine and non-routine displaced.

In the bottom panel, I test for heterogeneity depending on whether the workers are displaced in blue-collar occupations (service and sales, crafts, operators and assemblers, elementary occupations) or in white-collar occupations (managers, professionals, technicians and clerks).[8] Routine penalties among blue-collar workers are similar to those found for the entire sample. On the other hand, I find no evidence of routine penalties for white-collar displaced workers. This indicates that routine cognitive workers are able to cope with layoffs better than routine manual workers. The mechanisms behind this would be an interesting topic for further study.

[6] Only a quarter of the displaced workers have more than high school education and routine workers are underrepresented within this group. However, placing the threshold at a lower education level is problematic due to the changes to the Swedish primary and secondary schooling systems that affected different cohorts of workers.
[7] Including primary industries.
[8] Among blue-collar workers, routine three-digit occupations (according to the main definition) are found in the broad groups of crafts, operators and assemblers and elementary occupations. Among white-collar workers, routine three-digit occupations are found among clerks.

Figure 7. *Heterogeneity in routineness penalty in terms of post-displacement earnings by education level, industry and occupational group*

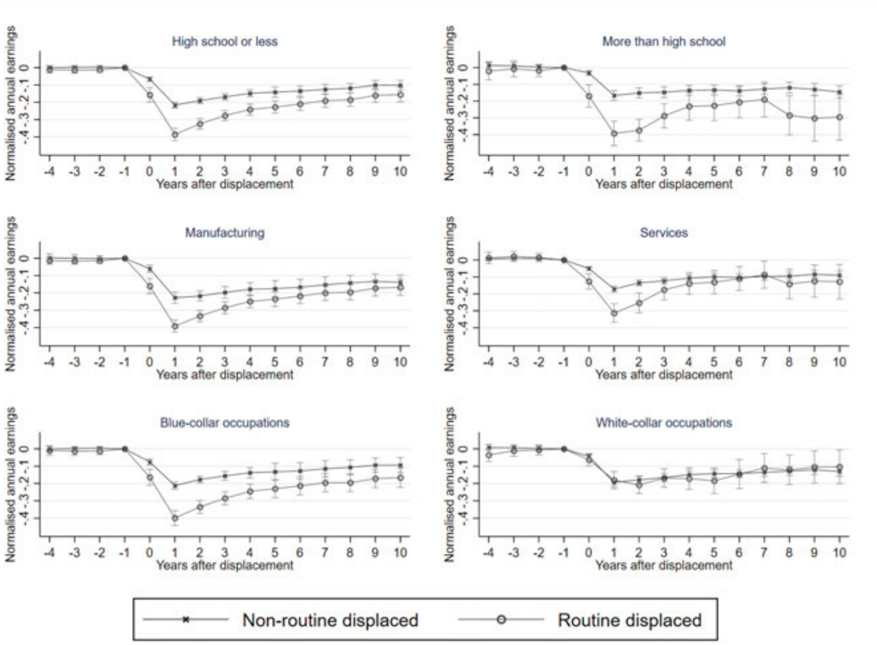

Note: The baseline period is the year before displacement, t_{-1}. Outcome of routine and non-routine displaced within the high/low educational groups, manufacturing/service industries and blue-collar/white-collar occupations relative to routine and non-routine matched controls in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

4.4 Mechanisms

Since routine occupations have been declining as a share of total employment, theory predicts that displaced routine workers should have a hard time finding a new job in their old occupation and may have to switch to another one when re-entering employment. This may lead both to adverse consequences in terms of earnings and wages relative to displaced non-routine workers in the short run as occupation-specific human capital is lost and to better long-run outcomes relative to routine workers who are not displaced and stay in declining occupations (Cortes, 2016). As routine occupations are concentrated in declining industries like manufacturing, displaced routine workers should be more likely to switch industry as well. The effects of industry switching are predicted to be qualitatively similar to those of occupation switching.

The probabilities that displaced routine and non-routine workers switch three-digit occupation and three-digit industry compared to the t_{-1} period are shown in Figures 8 and 9 respectively. Probabilities are conditional on the workers being employed in the given period; the occupational outcome is known only for a subset of employed workers, sampled according to the description in Section 2.[9] The results show that routine workers are more likely to change both occupation and industry in the years following displacement. For occupations, this effect is estimated to be 16 percentage points two years after displacement, when it peaks. It decreases somewhat over time, as non-displaced routine workers also switch occupations to a slightly higher degree, but remains at a significant nine percentage points in t_{10}. In the case of industries, routine workers are 20 percentage points less likely to be employed in their original industry than non-routine workers in year t_1. The gap remains at this level for the duration of the period over which the workers are followed. The results are qualitatively unaffected if a fully balanced panel consisting only of workers who are observed in each of the years t_{-4} to t_{10} is used, as can be seen in the Appendix Figure A9.

Figure 8. *Effect of displacement on the probability of routine and non-routine workers being in another three-digit occupation than in t_{-1}, conditional on being employed and occupational data being available*

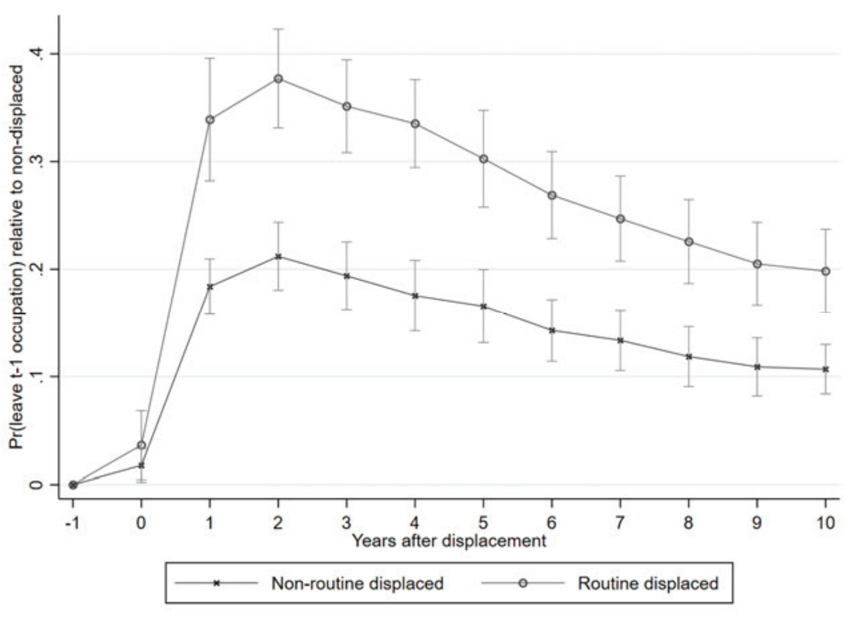

Note: Outcome of routine and non-routine displaced relative to routine and non-routine matched controls in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

[9] If occupation data is missing for an employed worker in the post-period, it is imputed according to the same procedure as is followed for t_{-1} occupations, and described in Section 2. If the occupation is still unknown post-imputation, the individual is dropped from the occupation-switching regression.

Figure 9. *Effect of displacement on the probability of routine and non-routine workers being in another three-digit industry than in the t_(-1) period, conditional on being employed*

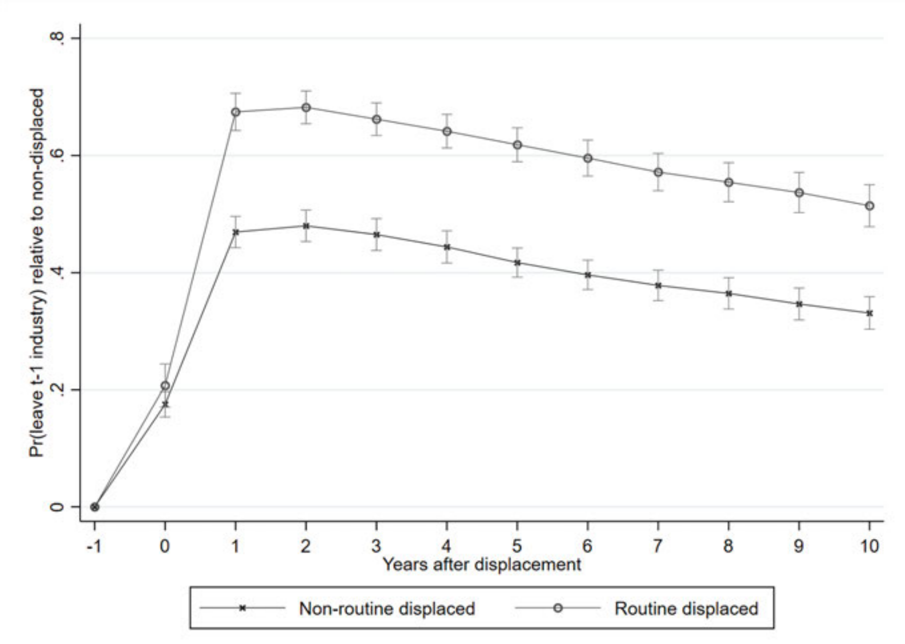

Note: Outcome of routine and non-routine displaced relative to routine and non-routine matched controls in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

Both the occupation and industry-switching results point to routine workers being more likely to become re-employed in jobs less similar to their pre-displacement jobs. Loss of occupation- and industry-specific human capital could therefore explain at least part of the additional short-run penalties that they suffer compared to their non-routine counterparts. However, higher rates of occupation and industry switching could also explain the faster long-term convergence of displaced routine workers' outcomes to those of their non-displaced peers. This is the case if their new occupations and industries see higher rates of wage growth than their original routine jobs; Cortes (2016) shows that routine workers who switch to non-routine jobs fare better over long time horizons than those who stay in routine occupations.

Figure 10. *Wage levels in displaced routine and non-routine workers' occupations and industries (conditional on being employed and occupational information available)*

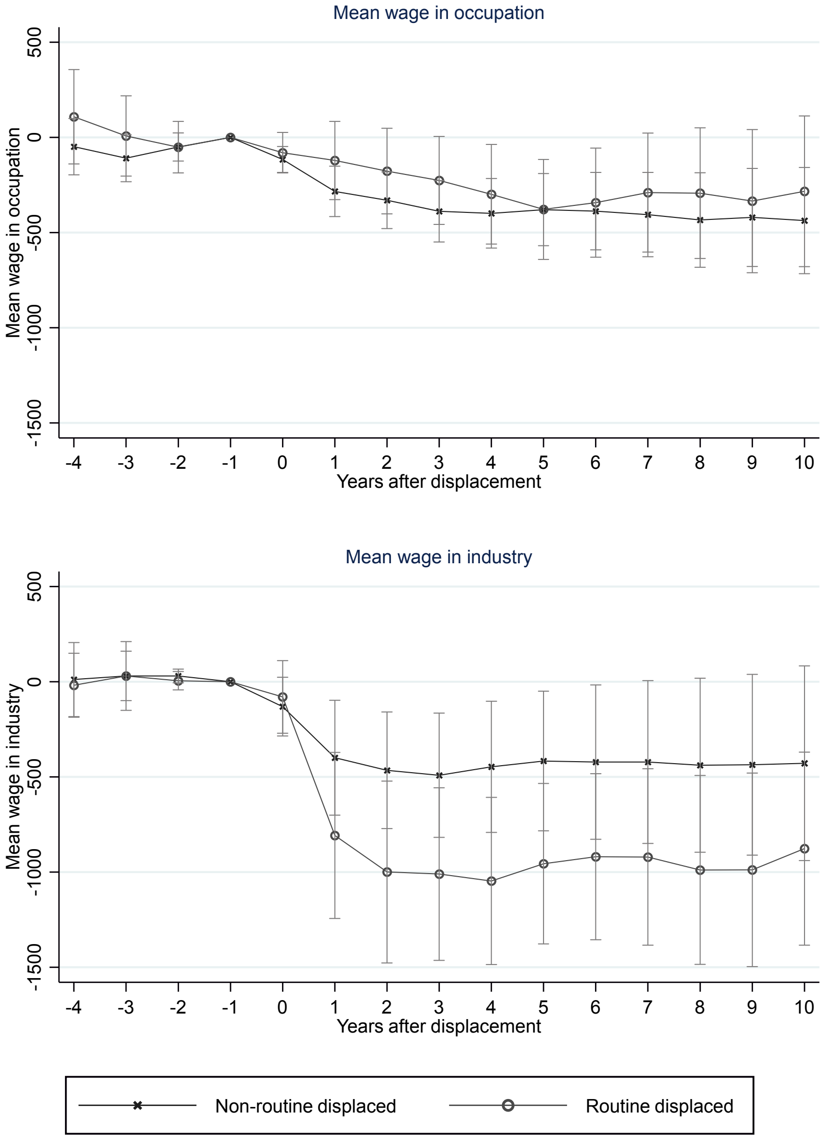

Note: Outcome of routine and non-routine displaced relative to routine and non-routine matched controls in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

Figure 11. *Wage levels of displaced routine and non-routine workers as share of the occupation or industry mean (conditional on being employed and occupational information available)*

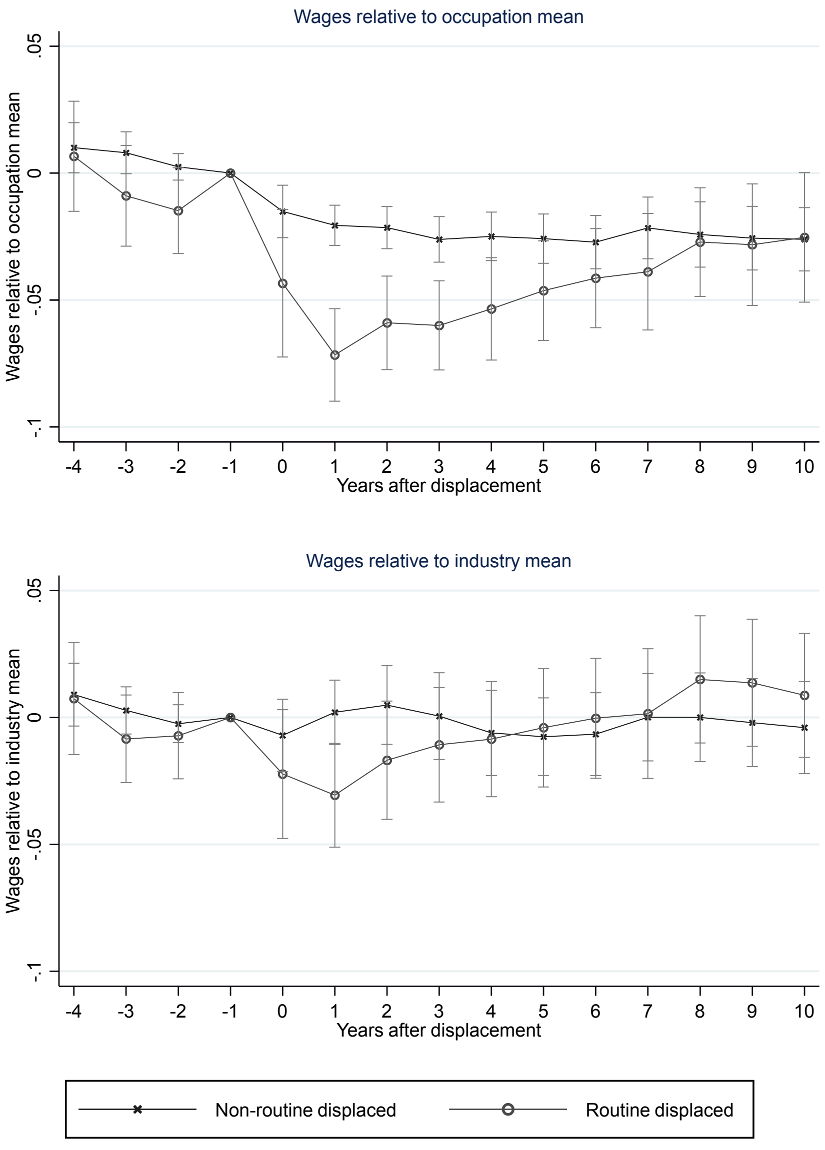

Note: Outcome of routine and non-routine displaced relative to routine and non-routine matched controls in each period. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

Does the higher switching rate among routine workers mean that they are more likely to move into high-paying occupations and industries? I analyse this in

Figure 10, which plots mean wages in the displaced individual's occupation and industry, relative to t_{-1}. This is based on the occupations and industries of both stayers and switchers, with the restriction that the individual must be employed and have available occupation and industry information respectively. The top panel shows that both routine and non-routine displaced workers tend to end up in lower-paying occupations than they were in in t_{-1}. The estimates are imprecise however, and there is no evidence that routine workers end up in lower-paying occupations than non-routine ones. When it comes to industry, there is a clearer pattern of routine workers ending up in lower-paying industries than the one that they were displaced from. This could be seen as evidence that they lose good industry matches or industry-specific rents. Many routine workers are displaced from manufacturing industries, which tend to provide high wages for less-educated individuals. Importantly, neither group of displaced workers manages to move up to better-paid sectors.

Is there evidence that routine workers move into jobs for which their human capital is less suited? I test this by considering the worker's wages as a percentage of the mean for their occupation and industry. The results are shown in Figure 11. Routine workers earn about seven percentage points less in terms of their occupation's mean in t_1, compared to about two percentage points for non-routine workers (relative to what had been the case in t_{-1}). The difference between the two groups remains significant until t_4, but the estimates for routine and non-routine workers converge in the long run. This could be due to routine workers having to enter occupations which differ more from their original one, resulting in a period of more rapid human capital accumulation. There are no clear patterns when it comes to wages relative to the industry mean.

As a final test of how occupational switching affects workers, the outcomes of displaced workers who are in a different occupation in t_5 are compared to those of displaced workers who remain in their t_{-1} occupation. This exercise is clearly endogenous to factors such as worker skill, motivation and local labour market conditions and is limited by the fact that occupations are observed for only a fraction of employed workers in the post-period. Nevertheless, it yields interesting indicative results as shown in Table 2, where the outcome is average relative earnings over the t_1 - t_5 period. Almost four fifths of the employed displaced routine workers with available occupation data were in another three-digit occupation in t_5; a full 70 percent of these switchers had gone to a non-routine occupation. Among non-routine workers, only six out of ten had switched out of their initial line of work. Of these, 92 percent went to another non-routine occupation. There is no evidence that switching occupations leads to higher earnings over the years t_1 - t_5. On the contrary, switching seems to be especially detrimental for routine workers, as they expect to lose seven percent of pre-displacement income annually if they switch to another routine occupation and 18 percent if they switch to a non-routine occupation. The losses for non-routine workers are smaller, at two percentage points if they are in another non-routine occupation and nine percentage points if they are in a routine occupation. Confidence intervals are tight and the estimates for stayers and switchers within each category are distinguishable at conventional significance levels. The results are evidence of loss of occupation-specific human capital hurting all workers, but especially routine ones. The better long-term prospects of non-routine occupations do not seem to help routine workers who switch into them in the short and medium run. These workers instead appear to lose more than those who switch to other routine occupations, which are more similar to the pre-displacement occupation in terms of tasks.

Table 2. *Average annual earnings of displaced workers in t_1-t_5 relative to the $t_{-4}-t_{-1}$ period depending on initial occupation routineness, whether the worker stayed in their initial occupation and the routineness of the new occupation conditional on switching*

	Stayed in routine occupation	Stayed in non-routine occupation	Switched to (other) routine occupation	Switched to (other) non-routine occupation
Routine occupation initially	1.07 (0.006) *N*=1,569		1.01 (0.007) *N*=1,742	0.90 (0.006) *N*=3,991
Non-routine occupation initially		1.14 (0.003) *N*=9,890	1.06 (0.010) *N*=1,241	1.12 (0.003) *N*=13,536

Note: Workers must be employed in t_5, with occupational information available, to be included in the analysis. Standard errors clustered at the level of t_{-1} establishments.

Further evidence on how earnings develop over time depending on displaced workers' occupation in t_5 is provided by the plots in Figure 12. To be included, workers must have an observed occupation in t_5. By definition, this means that they are employed in that period, which means that they are positively selected among displaced workers. Because of this, the comparison group of non-displaced workers is also limited to individuals whose occupations were observed in t_5. As occupation switching is endogenous to displacement, all three groups of displaced workers are compared to the entire sample of non-displaced workers, which is not split according to t_5 occupation. The caveat of the groups being non-randomly selected among displaced individuals still applies, but there nevertheless are interesting suggestive results. Both routine and non-routine workers seem to suffer in the short run if they switch occupation, but there are indications that each group suffers more if they

Figure 12. *Development of earnings for displaced workers over time, for occupation stayers, occupation switchers to non-routine occupations and occupation switchers to routine occupations (as defined by the occupation of individuals in t_5)*

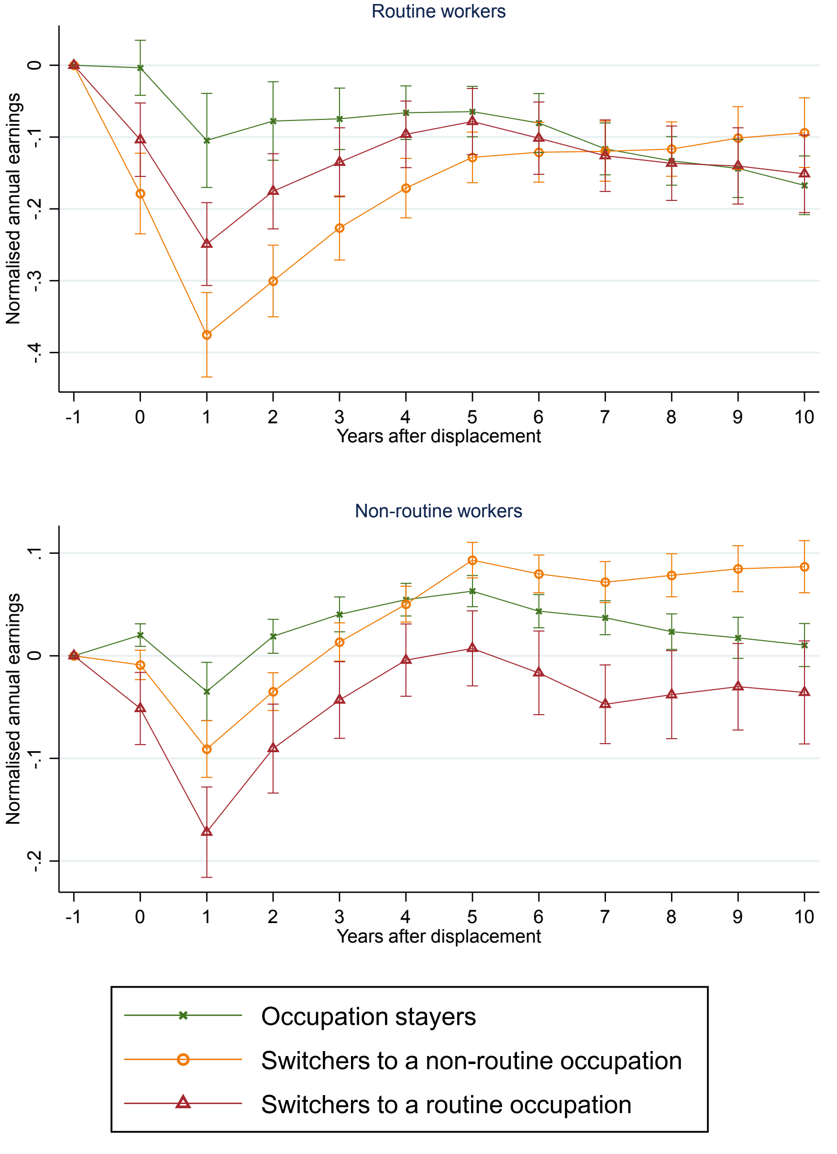

Note: Workers must be employed in t_5, with occupational information available, to be included in the analysis. Outcomes compared to those of controls which were employed and for whom occupational data were available in t_5. Standard errors clustered at the level of t_{-1} establishments. 95 percent confidence intervals shown.

switch to the other type of occupation. This is expected for non-routine workers if they switch to a routine occupation with bad prospects, but expectations are not quite as clear for routine workers who switch to non-routine occupations. That switchers do worse than stayers should be seen as a piece of evidence favouring the hypothesis that losses of occupation-specific human capital are an important component of displacement losses. Such losses should be larger if the worker switches to a more dissimilar occupation, which is what is indicated by the results. Occupation switchers, especially those who go to (other) non-routine occupations, do appear to gain on stayers over time. However, for routine workers, such gains are at most small and appear only towards the end of the period studied.

5. Conclusion

While large bodies of literature have identified that routine occupations have declined due to technological change and that workers lose out greatly in terms of their labour market outcomes following involuntary job loss, research connecting these two strands has been lacking. This paper attempts to conjoin the two lines of inquiry by comparing how workers initially in routine and non-routine occupations fare on the labour market following layoff. The findings imply substantial earnings, employment, wage and unemployment penalties of displacement for routine workers, up to several times the size of the penalties faced by non-routine displaced workers. These differences in losses persist in at least the medium run. There are indications that the additional losses suffered by routine workers are due to them being unable to find new jobs which provide a good match for their occupation- and industry-specific human capital, as they switch occupations and industries to a higher degree than displaced non-routine workers. This is reflected in routine workers moving to lower-paid industries and ending up on lower rungs in their occupations' wage distributions. Occupation switchers appear to do worse in terms of earnings than stayers, even if they switch to non-routine occupations. This is a somewhat disheartening piece of evidence for policy, which often aims to make displaced workers more flexible in terms of their job search and to re-educate and retrain them so that they can shift out of declining occupations and industries. A potential interpretation is that retraining programmes are insufficiently focused on the needs of displaced routine workers. Singling out this group as a target for such efforts and tailoring suitable courses might be a possible remedy.

References

Acemoglu, Daron and Autor, David. Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings in *Handbook of Labor Economics*, Volume 4, Part B, 2011

Autor, David H., Levy, Frank and Murnane, Richard J. The Skill Content of Recent Technological Change: An Empirical Exploration, *The Quarterly Journal of Economics*, Volume 118.4, 2003

Blien, Uwe, Dauth, Wolfgang and Roth, Duncan. Occupational Routine Intensity and the Costs of Job Loss: Evidence from Mass Layoffs. *Labour Economics*, Volume 68, 2021

Cortes, Guido Matias. Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data. *Journal of Labor Economics*, Volume 34.1, 2016

Cortes, Guido Matias, Jaimovich, Nir and Siu, Henry E. Disappearing routine jobs: Who, how, and why?, *Journal of Monetary Economics*, Volume 91, 2017

Cortes, Guido Matias, and Gallipoli, Giovanni. The costs of occupational mobility: An aggregate analysis. *Journal of the European Economic Association*, Volume 16.2, 2018

Dauth, Wolfgang, Findeisen, Sebastian and Suedekum, Jens. Adjusting to globalization in Germany. *Journal of Labor Economics*, Volume 39.1, 2021

Davis, Steven J. and von Wachter, Till M. *Recessions and the Cost of Job Loss*. NBER Working Paper No. 17638, Revised 2017

Eliason, Martin and Storrie, Donald. Lasting or Latent Scars? Swedish Evidence on the Long- Term Effects of Job Displacement. *Journal of Labor Economics*, Volume 24.4, 2006

Galaasen, Sigurd Mølster, and Kostøl, Andreas. *Mismatch and the Consequences of Job Loss*. WP, 2018

Goos, Maarten, Manning, Alan and Salomons, Anna. Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*, Volume 104.8, 2014

Graetz, Georg and Michaels, Guy. Robots at Work. *Review of Economics and Statistics*, Volume 100.5, 2018

Hethey-Maier, Tanja, and Schmieder, Johannes F. Does the use of worker flows improve the analysis of establishment turnover? Evidence from German administrative data. *NBER Working Paper No. w19730*, 2013.

Jaimovich, Nir, and Siu, Henry E. Job Polarization and Jobless Recoveries. *NBER Working Paper No. 18334*, 2018

Krolikowski, Pawel. Choosing a control group for displaced workers. *ILR Review*, Volume 71.5, 2018

Neffke, Frank, Nedelkoska, Ljubica and Wiederhold, Simon. Skill Mismatch and the Costs of Job Displacement. *CESifo Working Paper No. 9703*, 2022

Robinson, Chris. Occupational mobility, occupation distance, and specific human capital. *Journal of Human Resources*, Volume 53.2, 2018

Sullivan, Daniel, and Von Wachter, Till. Job displacement and mortality: An analysis using administrative data. *The Quarterly Journal of Economics*, Volume 124.3, 2009

Appendix

Table A1. *Pre-layoff levels of outcomes not used in matching*

	Controls (Unmatched)	**Displaced** (Unmatched)	**Controls** (Matched)	**Displaced** (Matched)
N individuals	1,035,499	84,896	65,069	84,325
Pre-period employment (probability)				
t_{-1}	1	1	1	1
t_{-2}	1	1	1	1
t_{-3}	1	1	1	1
t_{-4}	1	1	1	1
Pre-period log monthly wages				
t_{-1}	10.2	10.2	10.2	10.2
t_{-2}	10.2	10.1	10.2	10.1
t_{-3}	10.1	10.1	10.1	10.1
t_{-4}	10.1	10.1	10.1	10.1
Pre-period days of unemployment				
t_{-1}	1.7	9.7	2.5	9.7
t_{-2}	1.0	1.7	1.4	1.7
t_{-3}	1.7	2.5	2.5	2.5
t_{-4}	3.2	4.9	4.5	4.8

Note: Characteristics evaluated in year t_{-1} unless stated otherwise. Unmatched control group consists of 5% random sample of the eligible control pool. One-to-one propensity score matching with replacement implemented based on characteristics listed in the table. Propensity scores estimated using logit. Matched control sample statistics weighted by the number of times a control worker was drawn as the best match for a displaced worker. Sum of matched control weights is 84,325.

Table A2. *Descriptive statistics for the matched routine and non-routine samples of controls and displaced*

	Non-Routine (Matched controls)	**Non-Routine** (Matched displaced)	**Routine** (Matched controls)	**Routine** (Matched displaced)
N individuals	50,802	64,084	14,267	20,241
Routine intensity	0.50	0.50	0.81	0.81
Year t_{-1}	2004.3	2004.3	2004.3	2004.3
Age	43.5	43.4	42.8	43.1
Tenure	5.3	5.4	6.3	6.4
Female	0.38	0.38	0.34	0.32
Immigrant	0.10	0.11	0.14	0.15

Education level (percentages)				
Less than compulsory	5.72	5.41	11.92	12.01
Compulsory, 9 years	10.78	10.36	20.53	21.00
High school, 2 years	30.88	30.17	40.31	40.73
High school, 3 years	21.29	22.30	20.73	19.96
Some post-secondary	14.42	15.23	4.94	4.86
University	15.84	15.52	1.54	1.41
PhD	1.05	1.02	0.03	0.02
Occupations (percentages)				
Officials & Managers	8.94	9.04	0.00	0.00
Professionals	16.24	16.41	0.00	0.00
Technicians	22.63	22.86	0.00	0.00
Clerks	12.22	13.14	8.93	6.13
Service & Sales	13.18	13.31	0.00	0.00
Crafts	14.15	14.42	2.01	1.57
Operators & Assemblers	7.36	5.74	78.18	80.98
Elementary Occupations	5.27	5.07	10.88	11.33
Industries (percentages)				
Primary	0.59	0.50	0.05	0.08
Manufacturing	36.31	35.72	87.42	90.54
Construction	2.84	2.83	0.14	0.26
Utilities & telecom	11.67	12.11	3.99	2.00
Wholesale & retail	13.19	13.05	2.31	1.77
Business services	21.08	21.61	3.15	2.52
Health, social work	9.19	9.13	2.31	2.47
Education	2.08	1.95	0.07	0.09
Public administration	3.05	3.10	0.56	0.27
Type of municipality (percentages)				
Rural municipalities	12.82	11.68	26.00	28.27
Commuter municipalities	5.29	4.91	6.79	7.42
Towns	13.32	13.22	20.06	20.83
Other cities	32.07	33.60	31.22	29.67
Suburbs of 3 largest cities	11.81	11.50	5.87	5.48
3 largest cities	24.69	25.09	10.06	8.33

Pre-period earnings (SEK thousands)				
t_{-1}	323	334	333	334
t_{-2}	316	321	323	321
t_{-3}	306	308	311	309
t_{-4}	293	294	296	295

Note: Characteristics evaluated in year t_{-1} unless stated otherwise. Workers subdivided according to the main definition of routineness, as defined in Section 2.1. Matched control sample statistics weighted by the number of times a control worker was drawn as the best match for a displaced worker.

Estimates of annual earnings in different periods for different worker groups are shown in Table A3. The first set of estimates presents the period effects relative to t_{-1} for non-routine non-displaced workers. The second set contains interactions of periods with routineness, and provides estimates of the difference between routine and non-routine workers' relative earnings in each period. The third set contains interactions of each period with displacement; these estimates show the difference between displaced and non-displaced non-routine workers. Finally, the fourth set of estimates is the focus of this paper, as it shows the difference between routine and non-routine displaced workers. The total size of the displacement effect for routine workers is found by adding the period-displacement and the period-displacement-routine effect for the period in question.

The main definition of routineness used in this paper (the quarter of workers with the highest share of routine tasks in their occupations classified as routine, the others as non-routine) is used in columns (1) and (2). In columns (3) and (4), the definition of routine occupations is made less stringent and the median routineness of workers' t_{-1} occupations is used as the cutoff. The final two columns provide estimates in the case when only the top and bottom quartile of routineness are included, so as to ensure that truly routine workers are compared to truly non-routine ones. For each definition of routineness, estimates are provided for the full unmatched sample as well as for a sample that has been matched on covariates and earnings pre-trends as described in Section 4. The main specification used in this study is the one in column (2).

Table A3. *Point estimates of period effects for different groups of workers*

	(1) Baseline, un-matched	(2) Baseline, matched	(3) Median cutoff, un-matched	(4) Median cutoff, matched	(5) High and low only, un-matched	(6) High and low only, matched
Periods:						
t_{-4}	-0.11***	-0.13***	-0.11***	-0.14***	-0.12***	-0.14***
	(0.0006)	(0.003)	(0.0006)	(0.004)	(0.0009)	(0.005)
t_{-3}	-0.063***	-0.079***	-0.063***	-0.082***	-0.069***	-0.091***
	(0.0005)	(0.003)	(0.0005)	(0.003)	(0.0007)	(0.004)
t_{-2}	-0.025***	-0.036***	-0.025***	-0.039***	-0.028***	-0.046***
	(0.0003)	(0.002)	(0.0003)	(0.002)	(0.0005)	(0.002)
t_0	0.0052***	-0.000073	0.0055***	0.0025	0.011***	0.0092***
	(0.0004)	(0.002)	(0.0004)	(0.002)	(0.0006)	(0.003)

t_1	0.013***	0.0073**	0.014***	0.014***	0.022***	0.020***
	(0.0006)	(0.003)	(0.0007)	(0.003)	(0.0009)	(0.004)
t_2	0.029***	0.021***	0.030***	0.030***	0.039***	0.033***
	(0.0009)	(0.004)	(0.0009)	(0.004)	(0.001)	(0.005)
t_3	0.048***	0.043***	0.050***	0.053***	0.059***	0.057***
	(0.001)	(0.004)	(0.001)	(0.005)	(0.002)	(0.006)
t_4	0.068***	0.061***	0.071***	0.074***	0.083***	0.081***
	(0.001)	(0.005)	(0.001)	(0.005)	(0.002)	(0.007)
t_5	0.091***	0.081***	0.094***	0.099***	0.11***	0.10***
	(0.002)	(0.006)	(0.002)	(0.006)	(0.003)	(0.008)
t_6	0.12***	0.10***	0.12***	0.12***	0.13***	0.13***
	(0.002)	(0.007)	(0.002)	(0.007)	(0.003)	(0.009)
t_7	0.14***	0.12***	0.14***	0.14***	0.16***	0.15***
	(0.002)	(0.008)	(0.002)	(0.008)	(0.003)	(0.01)
t_8	0.17***	0.15***	0.17***	0.17***	0.18***	0.18***
	(0.003)	(0.009)	(0.003)	(0.009)	(0.004)	(0.01)
t_9	0.19***	0.17***	0.20***	0.19***	0.21***	0.20***
	(0.003)	(0.010)	(0.003)	(0.010)	(0.004)	(0.01)
t_{10}	0.22***	0.20***	0.23***	0.22***	0.24***	0.23***
	(0.003)	(0.01)	(0.003)	(0.01)	(0.005)	(0.01)
Period-routine interactions:						
t_{-4}	0.022***	0.033***	0.012***	0.035***	0.030***	0.047***
	(0.002)	(0.006)	(0.001)	(0.005)	(0.002)	(0.007)
t_{-3}	0.016***	0.021***	0.0083***	0.019***	0.022***	0.034***
	(0.001)	(0.005)	(0.0010)	(0.004)	(0.001)	(0.006)
t_{-2}	0.0086***	0.0090**	0.0040***	0.011***	0.012***	0.019***
	(0.0007)	(0.003)	(0.0006)	(0.003)	(0.0008)	(0.004)
t_0	-0.011***	-0.017***	0.015*	-0.013***	-0.016***	-0.027***
	(0.001)	(0.004)	(0.007)	(0.003)	(0.001)	(0.005)

t_1	-0.022***	-0.042***	0.0061	-0.034***	-0.031***	-0.054***
	(0.001)	(0.006)	(0.006)	(0.005)	(0.002)	(0.006)
t_2	-0.036***	-0.060***	0.0017	-0.047***	-0.046***	-0.072***
	(0.002)	(0.007)	(0.005)	(0.006)	(0.002)	(0.008)
t_3	-0.050***	-0.075***	-0.0052***	-0.057***	-0.062***	-0.089***
	(0.002)	(0.009)	(0.0007)	(0.007)	(0.003)	(0.009)
t_4	-0.061***	-0.087***	-0.011***	-0.070***	-0.076***	-0.11***
	(0.003)	(0.01)	(0.001)	(0.009)	(0.003)	(0.01)
t_5	-0.073***	-0.094***	-0.017***	-0.082***	-0.088***	-0.12***
	(0.003)	(0.01)	(0.002)	(0.010)	(0.004)	(0.01)
t_6	-0.083***	-0.11***	-0.025***	-0.092***	-0.10***	-0.13***
	(0.004)	(0.01)	(0.002)	(0.01)	(0.005)	(0.01)
t_7	-0.095***	-0.12***	-0.032***	-0.095***	-0.11***	-0.15***
	(0.005)	(0.02)	(0.002)	(0.01)	(0.005)	(0.02)
t_8	-0.11***	-0.13***	-0.038***	-0.10***	-0.12***	-0.16***
	(0.005)	(0.02)	(0.003)	(0.01)	(0.006)	(0.02)
t_9	-0.12***	-0.14***	-0.044***	-0.11***	-0.13***	-0.17***
	(0.006)	(0.02)	(0.004)	(0.02)	(0.007)	(0.02)
t_{10}	-0.13***	-0.15***	-0.051***	-0.12***	-0.15***	-0.18***
	(0.006)	(0.02)	(0.004)	(0.02)	(0.008)	(0.02)
Period-displaced interactions:						
t_{-4}	-0.018**	0.0041	-0.026***	0.0047	-0.027***	0.0069
	(0.007)	(0.007)	(0.008)	(0.008)	(0.005)	(0.006)
t_{-3}	-0.012*	0.0048	-0.019**	0.0031	-0.018***	0.0091
	(0.006)	(0.006)	(0.007)	(0.007)	(0.004)	(0.005)
t_{-2}	-0.0096*	0.0033	-0.014*	0.0017	-0.010***	0.0084*
	(0.004)	(0.005)	(0.006)	(0.006)	(0.003)	(0.004)
t_0	-0.059***	-0.056***	-0.050***	-0.049***	-0.040***	-0.039***
	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)	(0.006)

t_1	-0.20***	-0.20***	-0.19***	-0.19***	-0.16***	-0.16***
	(0.009)	(0.009)	(0.01)	(0.01)	(0.01)	(0.01)
t_2	-0.18***	-0.18***	-0.17***	-0.17***	-0.15***	-0.15***
	(0.009)	(0.009)	(0.01)	(0.01)	(0.01)	(0.01)
t_3	-0.16***	-0.16***	-0.15***	-0.16***	-0.14***	-0.14***
	(0.010)	(0.010)	(0.01)	(0.01)	(0.01)	(0.01)
t_4	-0.15***	-0.14***	-0.14***	-0.14***	-0.14***	-0.13***
	(0.010)	(0.01)	(0.010)	(0.01)	(0.01)	(0.01)
t_5	-0.15***	-0.14***	-0.13***	-0.14***	-0.14***	-0.13***
	(0.01)	(0.01)	(0.009)	(0.01)	(0.01)	(0.01)
t_6	-0.15***	-0.14***	-0.13***	-0.14***	-0.14***	-0.13***
	(0.01)	(0.01)	(0.010)	(0.01)	(0.01)	(0.01)
t_7	-0.14***	-0.13***	-0.12***	-0.12***	-0.13***	-0.12***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
t_8	-0.14***	-0.12***	-0.12***	-0.12***	-0.13***	-0.13***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
t_9	-0.13***	-0.11***	-0.12***	-0.11***	-0.13***	-0.12***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
t_{10}	-0.13***	-0.12***	-0.12***	-0.12***	-0.14***	-0.13***
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.02)
Period-displaced-routine interactions:						
t_{-4}	-0.012	-0.015	0.015*	-0.0096	-0.0029	-0.018
	(0.009)	(0.010)	(0.007)	(0.008)	(0.01)	(0.01)
t_{-3}	-0.018*	-0.018*	0.0061	-0.0060	-0.013	-0.023*
	(0.008)	(0.009)	(0.006)	(0.007)	(0.01)	(0.01)
t_{-2}	-0.017*	-0.016*	0.0017	-0.0051	-0.016	-0.021*
	(0.007)	(0.007)	(0.005)	(0.006)	(0.009)	(0.009)
t_0	-0.11***	-0.10***	-0.077***	-0.013***	-0.13***	-0.12***
	(0.02)	(0.02)	(0.01)	(0.003)	(0.02)	(0.02)

t_1	-0.20***	-0.19***	-0.14***	-0.034***	-0.24***	-0.22***
	(0.02)	(0.02)	(0.01)	(0.005)	(0.02)	(0.02)
t_2	-0.17***	-0.15***	-0.12***	-0.047***	-0.19***	-0.18***
	(0.02)	(0.02)	(0.01)	(0.006)	(0.02)	(0.02)
t_3	-0.13***	-0.12***	-0.10***	-0.057***	-0.15***	-0.14***
	(0.01)	(0.01)	(0.01)	(0.007)	(0.02)	(0.02)
t_4	-0.11***	-0.098***	-0.096***	-0.070***	-0.12***	-0.11***
	(0.02)	(0.02)	(0.01)	(0.009)	(0.02)	(0.02)
t_5	-0.094***	-0.090***	-0.094***	-0.082***	-0.11***	-0.098***
	(0.02)	(0.02)	(0.02)	(0.010)	(0.02)	(0.02)
t_6	-0.077***	-0.075***	-0.081***	-0.092***	-0.089***	-0.080***
	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
t_7	-0.071***	-0.065***	-0.081***	-0.095***	-0.081***	-0.070***
	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
t_8	-0.067***	-0.073***	-0.075***	-0.10***	-0.073***	-0.067**
	(0.02)	(0.02)	(0.02)	(0.01)	(0.02)	(0.02)
t_9	-0.058***	-0.060**	-0.059***	-0.11***	-0.058**	-0.051*
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
t_{10}	-0.044*	-0.050*	-0.052***	-0.12***	-0.039	-0.035
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)

Note: Standard errors clustered at the level of t_{-1} establishments shown in parentheses. * p < 0.05, ** p < 0.01, *** p < 0.001

Figure A1. *Histograms of propensity scores for the control and displaced samples before and after matching*

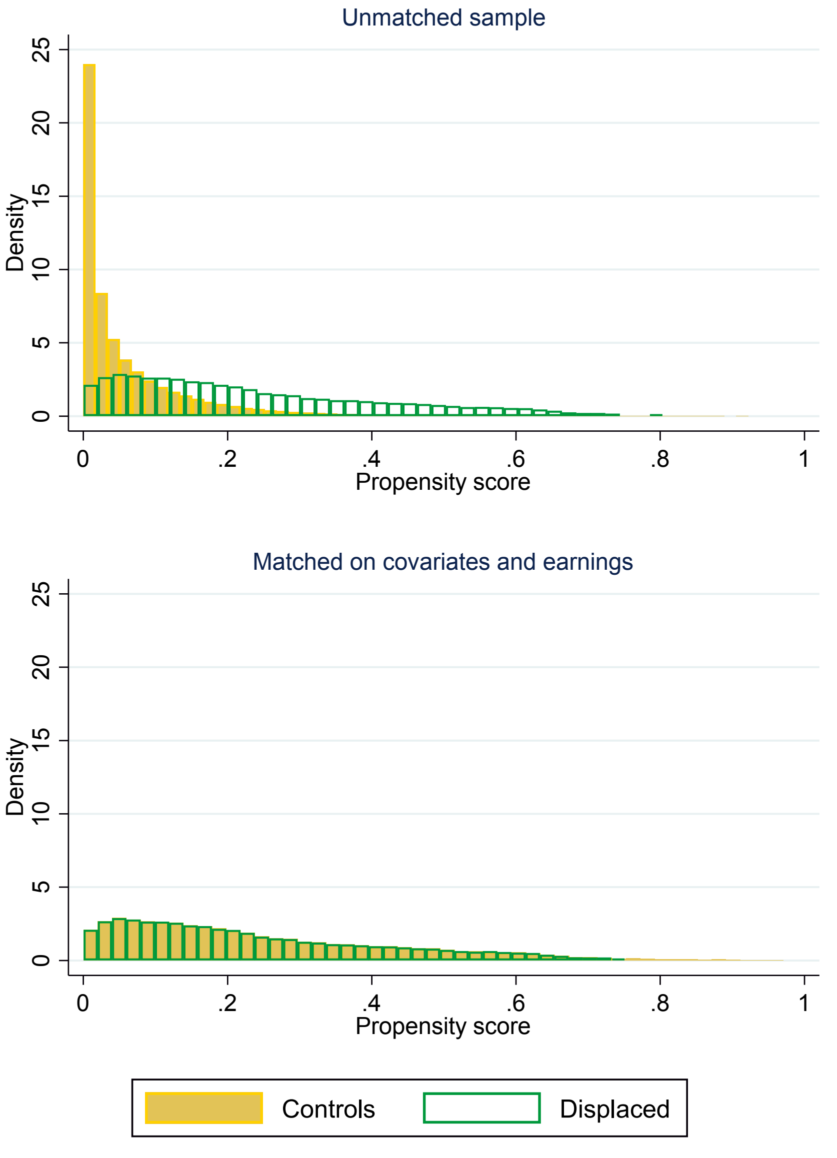

Figure A2. *Estimates of average displacement effects, without routine interactions, on employment, monthly wages and days of unemployment*

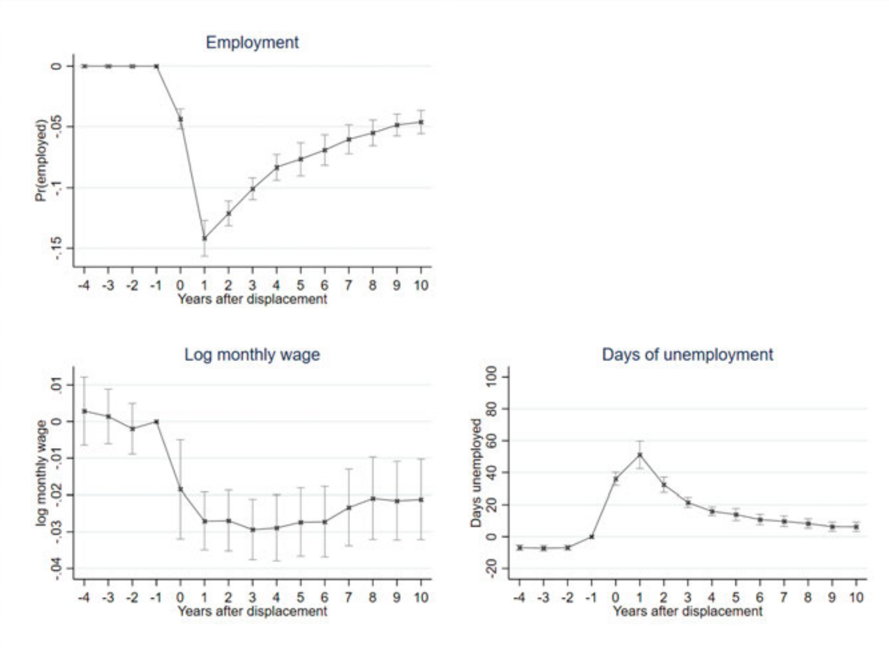

Figure A3. *Estimates of displacement earnings penalties for the four routineness quartiles*

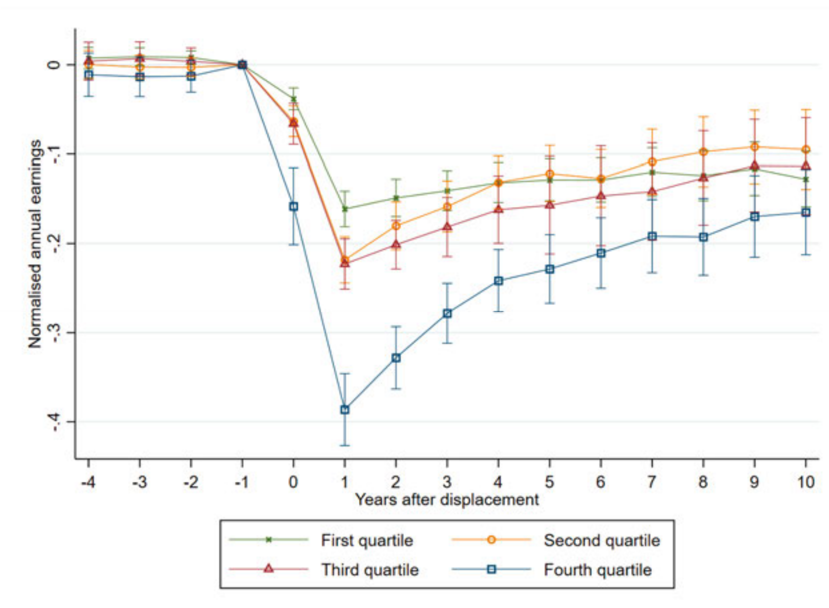

Figure A4. *Estimates of displacement employment penalties using different definitions of routine and non-routine occupations*

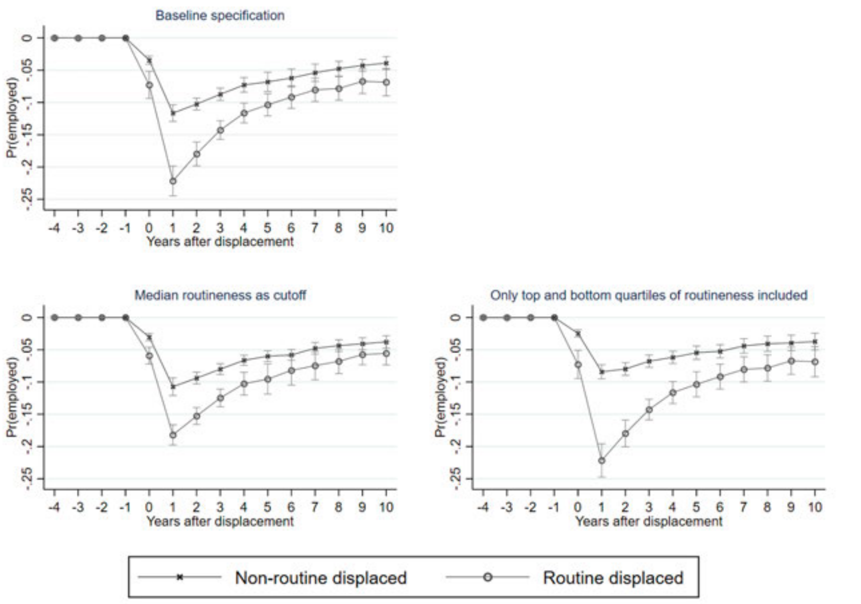

Figure A5. *Estimates of displacement monthly wage penalties using different definitions of routine and non-routine occupations*

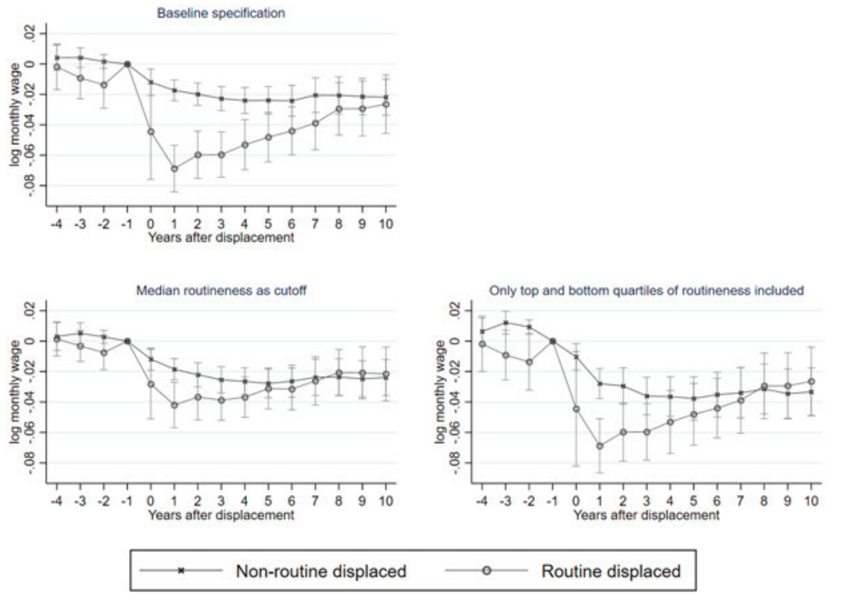

Figure A6. *Estimates of displacement unemployment penalties using different definitions of routine and non-routine occupations*

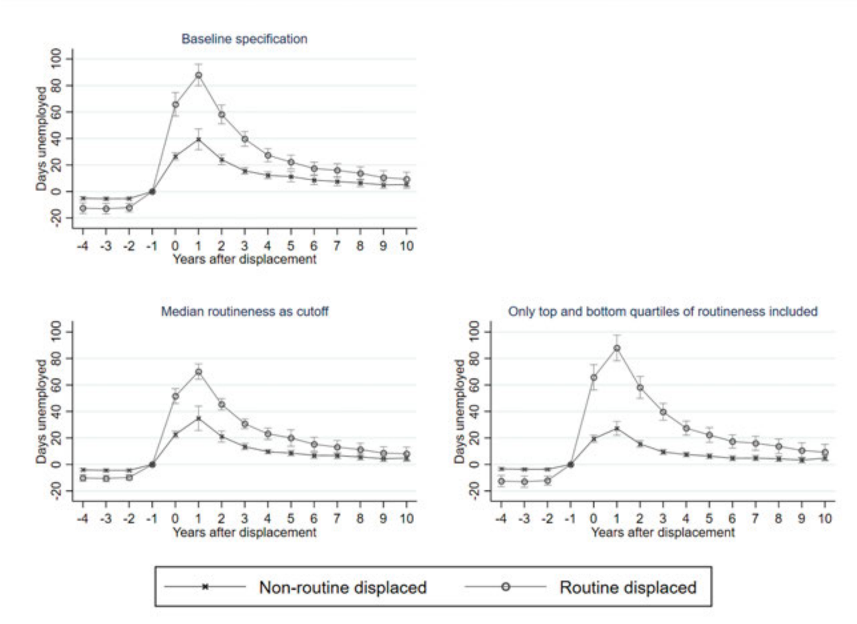

Figure A7. *Effects of displacement on annual earnings using the unmatched sample, a sample matched on a broad set of control variables and the baseline specification with a sample matched on a broad set of control variables and earnings pre-trends*

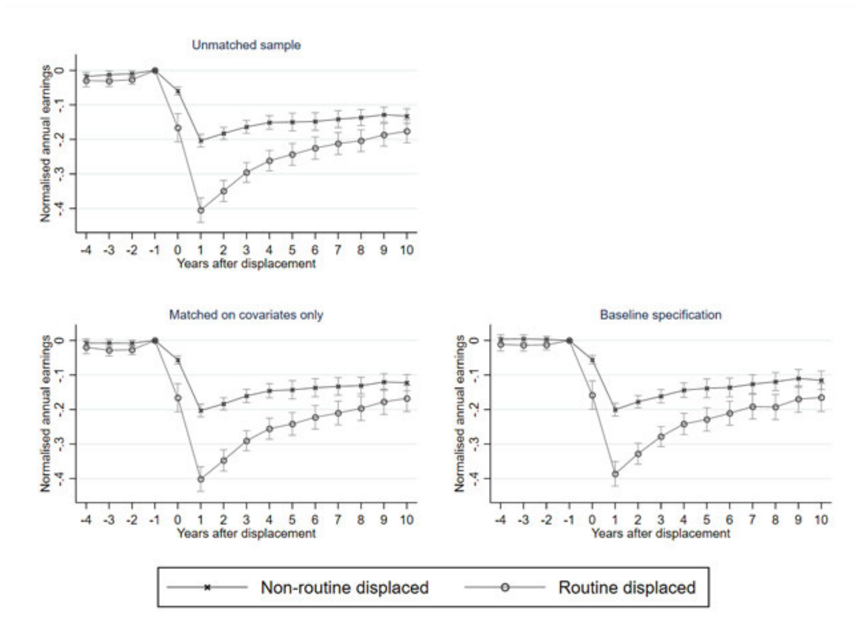

Figure A8. *Estimated effects of displacement on routine and non-routine workers using the full sample (left column) and only those individuals observed in each of the years t_{-4} to t_{10} (right column) on annual earnings and days in unemployment*

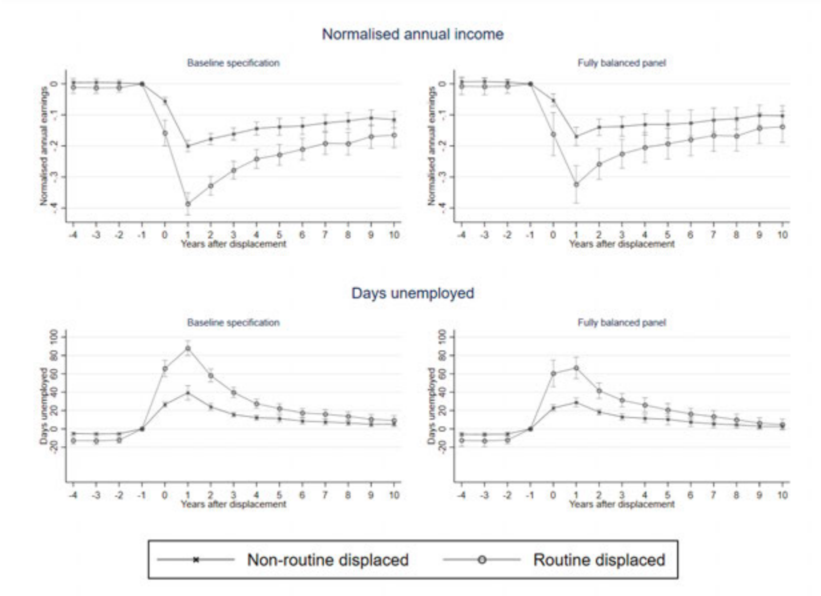

Figure A9. *Estimated effects of displacement on routine and non-routine workers using the full sample (left column) and only those individuals observed in each of the years t_{-4} to t_{10} (right column) on probability of switching occupation and switching industry*

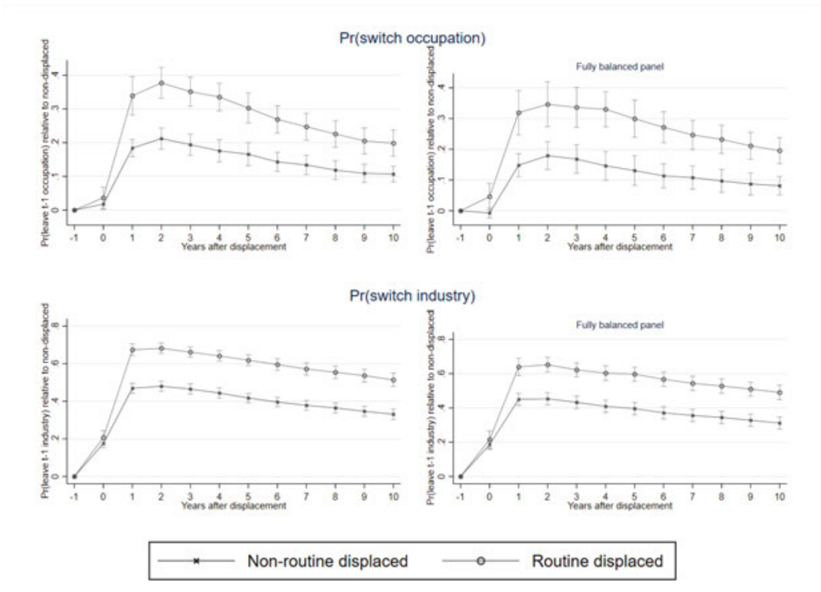

Essay III. Worker Attributes, Aggregate Conditions and the Impact of Adverse Labor Market Shocks

with Susan Athey, Lisa Simon, Oskar Nordström Skans and Johan Vikström

We thank Vitor Hadad, David Strömberg, Nicolaj Mühlbach, Stefan Pitschner, Stefan Eriksson and seminar participants at Stanford HAI and Uppsala University for their helpful suggestions and comments. Skans and Yakymovych gratefully acknowledge financial support from Vetenskapsrådet, grant number 2018-04581.

1. Introduction

The process of job reallocation across firms drives economic growth and creates benefits for society as a whole (Bartelsman and Doms, 2000). At the same time, the process displaces workers, who often suffer large and persistent earnings losses (Jacobson et al., 1993). These earnings losses can be sizeable enough to further impact the affected workers' health and well-being.[24] Consequently, governments across the OECD spend vast resources on an array of social programs designed to mitigate the earnings losses of displaced workers (OECD, 2019). The policy mix includes unemployment insurance or welfare benefit transfers, active labor market policies aimed at easing the transition between jobs, and policies such as employment protection legislation and short-time work schemes that aim to protect workers from being displaced under certain conditions. In the interest of targeting scarce public resources where they are needed the most, policymakers have to know which workers suffer the largest earnings losses. This paper's goal is to improve understanding in this area by identifying worker characteristics and aggregate conditions that predict the size of earnings losses after job displacement.

To this end, we use unusually rich administrative data to characterize individuals who lose their jobs in establishment closures, as well as the aggregate conditions prevailing at the time. We then estimate how such worker attributes and labor market characteristics interact to predict workers' resilience to job displacement. This is also informative about workers' ability to cope with adverse labor market shocks more generally. Using closures instead of individual unemployment spells allows us to study negative shocks which are well-identified in time, unrelated to other changes in personal circumstances, and well-identified even for workers who are displaced but manage to find new employment without an intermission. The last point is of specific importance, as we are interested in identifying conditions under which workers manage well when hit by adverse shocks. As those who lose their jobs in establishment closures are different from those who do not, we rely on propensity score matching to construct a comparable control group. Then, using selection-on-observables assumptions (unconfoundedness), we identify the counterfactual trajectories of displaced workers in the absence of a shock.

We proceed in the trail of a vast literature on the impact of mass layoffs and establishment closures pioneered by Jacobson et al. (1993). A large set of studies has illustrated heterogeneity in post-displacement outcomes due to factors on both the supply and demand sides of the labor market. On the supply side, much of the literature has emphasized the role of human capital, gender and age (see Davis and von Wachter, 2011, for an overview), but also attributes related to mobility, such as family status (Huttunen and Kellokumpu,

[24] See e.g. Eliason (2014a) on increased drinking, Black et al. (2015) on increased smoking and Eliason (2012) on increased divorces.

2016). On the demand side, we find a more scattered set of articles emphasizing the role of job content (Blien et al., 2021; Yakymovych, 2022), firm-specific wage levels (Lachowska et al., 2020; Gulyas and Pytka, 2019), the size of the displacement event (Gathmann et al., 2020; Cederlof¨, 2019), the sector (Eliason, 2014b) and aggregate business cycle conditions (Eliason and Storrie, 2006; Davis and von Wachter, 2011).

The Swedish administrative data used in this paper enable us to paint a uniquely rich picture of aspects that may affect the impact of job loss. We characterize displaced workers in terms of general and job-specific human capital (measured by years of schooling, experience and tenure), detailed family status, internal and external migration history, as well as a broad set of pre-displacement job characteristics capturing match quality, task content, establishment size, trends and wage premia. We measure aggregate conditions by characterizing the worker's pre-displacement industry and location. Industry characteristics include measures of turbulence and pre-existing and forward-looking trends, whereas locations are characterized by population density, unemployment rates and industry structure.

To understand how such a large number of aspects interact to determine the severity of displacement losses, we need to employ a method that allows estimation of multidimensional treatment effect heterogeneity in a meaningful and flexible way. Using our exceptionally rich data, we estimate a *Generalized Random Forest* (GRF; developed by Athey et al., 2019) on the set of displaced and matched control workers. The forest iterates across the dataset, splitting it based on the included covariates so as to maximize treatment effect heterogeneity across "leaves". It allows us to estimate a conditional average treatment effect (CATE) for each worker, based on the neighbours in his or her "leaf". This makes it possible to classify individuals based on how resilient they are to job displacement and to identify worker attributes and aggregate conditions that are associated with large displacement losses. GRF has several advantages relative to traditional sub-group analysis. It is able to consider heterogeneity across a large number of covariates at the same time. Furthermore, GRF employs a number of measures to minimize overfitting (replication of random patterns in the data), which would be a problem when including many covariates in a traditional heterogeneity analysis. GRF is also very flexible with regard to functional form and high-level interactions between variables.

Our short-run estimates suggest that displaced workers on average experience a 24 percent reduction in earnings and a 15 percentage point reduction in employment probability in the calendar year after the establishment closure. About one third of these effects still remain 10 years later. We also document a distinct and persistent impact on the probability of moving location and switching industry. There is substantial systematic heterogeneity around these averages across worker groups. Although we focus on heterogeneity in the short-run earnings effect, workers whom the GRF predicts to be more resilient in the short run also have substantially lower earnings and employment losses during the entire 10 year-long follow-up period. This is especially clear when considering the hardest-hit workers. In the year after displacement, workers in the quartile with largest predicted losses suffer a decrease in earnings of 40 percent (twice the median), and those in the hardest-hit decile suffer a massive earnings loss of almost 50 percent (2.5 times the median). This difference is persistent over time, and hardly decreases at all in relative terms. This suggests that our estimates capture fundamental differences in the resilience in the face of job loss.

Many attributes emphasized in the earlier literature are related to the heterogeneity in resilience to job loss that we observe. Workers in the hardest-hit group are older, less educated, have lower pre-displacement earnings, are laid off from establishments with a larger market share, are more often displaced from routine intensive jobs in manufacturing, and work in declining industries in rural areas. At the same time, the characterisation provided by the GRF makes it clear that none of these characteristics has a deterministic relationship with large displacement losses on its own. In particular, we focus on age and schooling, which are important predictors of displacement losses. To validate the GRF's predictions that older and less-educated workers lose more, we estimate differences in post-displacement earnings between displaced and controls within cells defined by combinations of age and schooling.[25] Earnings losses due to displacement are in excess of 50 percent for the oldest and least educated, and smaller than 10 percent for the youngest and most educated. We further let the GRF find the most resilient and least resilient quartiles of workers within each combination of age and schooling and estimate the earnings losses for these subsamples. For almost all combinations of age and schooling, the least resilient quartile of workers suffers earnings losses in excess of 30 percent, whereas the most resilient quartile of workers loses less than 20 percent. This suggests that there is considerable systematic heterogeneity which arises from other attributes even conditional on age and schooling. We show that much of this remaining heterogeneity can be explained by characteristics of industries (e.g., manufacturing) and of locations (e.g., population density).

As the results suggest that aggregate conditions at the industry and location level play an important role, we explore these features further. To make progress on this front, we need to deal with the fact that different features are correlated. For example, rural locations may have a high unemployment rate, and their residents may have lower schooling. We therefore use our GRF model to predict displacement effects for all workers in our sample as if they were displaced in each location or industry, but holding their individual and workplace characteristics constant. This gives us a measure of how severe the

[25] We often employ such displaced-control differences in our analysis, using the GRF to classify workers in terms of resilience and then validating the predictions in this way. Calculating displaced-control differences within groups defined by the GRF reduces issues related to model calibration. Throughout the paper, we use "ATE" as a synonym for such within-group displaced-control differences, while "CATE" refers to GRF estimates.

displacement effects would be in each location and industry for *the average displaced worker in the economy* as a function of the full set of location- or industry-specific attributes. Our results suggest that the model has been able to correctly identify “good” and “bad” industries and locations. The impact of displacement in “good” industries and locations is substantially lower than in “bad” industries and locations. “Good” locations have high population density and low unemployment rates, while “bad” ones are rural areas exposed to disadvantageous industry trends. “Bad” industries have negative long-run and short-run employment trends, low churning and reallocation rates, and are often found in manufacturing.[26] If displaced in ”good” industry- and location-level conditions, even workers who are old and have few years of schooling can cope fairly well.

Since job loss is likely to matter more for immobile workers, we also document how displaced workers move across locations and industries after displacement. Workers are more likely to move location and switch industry if they were displaced under bad conditions. Nevertheless, this increased mobility does not compensate them in terms of earnings. This might be because rates of geographical mobility for all groups of workers in our sample remain very low, in spite of the increase due to displacement.

We end the paper with a set of targeting exercises. Here, we illustrate how well a policymaker or forecaster could identify hard-hit workers if they split the sample on one or two easily observed attributes at a time compared to using GRF. We conclude that it is difficult to design a simple rule that can come close to being as good at identifying vulnerable workers as the full GRF model. The best simple rule for finding the hardest-hit group involves targeting older workers in manufacturing, but still underperforms relative to GRF. This points to the importance of a considerable number of interacting characteristics in determining the size of displacement losses.

Overall, we believe that the findings contribute important new insights to the literature on the impact of job displacement. The study most closely related to this paper is the one by Gulyas and Pytka (2019), who use GRF to study the impact of mass layoffs, but with a somewhat different focus (motivated by competing theories) and fewer included characteristics. Their results indicate that displacement losses are primarily related to firm-level rents before displacement. Our results do not corroborate this finding. On the contrary, they highlight the complexity of interacting factors that contribute to shaping the magnitudes and persistence of individual workers’ earnings losses. Age, schooling, industry and location all play a part.

The remainder of this paper is structured as follows. In Section 2 we give a condensed presentation of our data (details are in the appendix), as well as a

[26] Our model includes a manufacturing dummy, but GRF captures most of its impact if we exclude it and only use other industry characteristics.

description of the displaced workers, and illustrate the average impact of displacement. Section 3 outlines the estimation of the GRF, while Section 4 presents calibration exercises and describes the heterogeneity that we find. In Section 5, we discuss how different characteristics are related to the heterogeneous effects of displacement. Results related to aggregate industry- and location-level conditions are shown in Section 6. Section 7 reports results from the targeting exercises and Section 8 concludes.

2. Data Definitions, Matching and Main Effects of Displacement

2.1. Displaced Workers and Control Workers

The data source which defines our study population is the Swedish linked employer-employee register *RAMS*. These annual data files contain information on how much every establishment reimbursed each one of its employees during the year. An establishment is a production unit with a physical location within a firm (or other legal entity).[27] For simplicity, we will use the term ”firm” for all types of legal entities. The data are linked to various other records through person, establishment and firm identification numbers.

We consider displacement events in the years 1997-2014. By starting the analysis in 1997, we are able to use a pre-displacement period of 10 years to measure the labor market and mobility trajectories of workers leading up to the event. Ending in 2014 means that we can use at least 5 post-displacement years to measure outcomes for all displaced workers.

To define our sample, we start by selecting a panel of all individuals in Sweden aged 16 to 64. Individuals are defined as employed if they earn more than 3 times the minimum monthly wage from a single employer during the year.[28] For each year, we only retain an employed individual's main job, defined as the establishment from which he or she collected the largest amount of earnings during the year. An establishment's size is determined by the number of employed workers for whom it is the main place of work.

For every year t, we define *closing establishments* as those which had at least 5 employees in year t_{-1} and either $i)$ disappear completely by year t_1, or $ii)$ see the number of employees fall at least 90 percent from year t_{-1} to year

[27] Around 10 percent of workers are not employed at a physical establishment due to the nature of their work (e.g. home-care workers). These workers can therefore never be included in the set of displaced or the control group. If a worker transitions into employment without a physical establishment during the outcome years, we do include him or her in our measures of earnings and employment.

[28] Sweden does not have a legislated minimum wage. Following conventions in studies of Swedish register data, we instead proxy it by the 10th percentile of the wage distribution.

t_1. We further require that the establishment had some economic activity during year t_0. We remove *false closures*, defined as cases where at least 30 percent of workers involved in an apparent closure moved to a single new establishment, or to other establishments within the original firm. This is a standard procedure in the literature (see Kuhn, 2002, for a detailed discussion), as these are likely to be mergers, splits or reorganisations rather than actual establishment closures.[29]

Our group of *displaced workers* consists of those who were aged 24 to 60 and employed at a closing establishment in year t_{-1}. The age restrictions are imposed to ensure that we can measure pre-displacement characteristics and post-displacement outcomes in a meaningful way (for instance, to avoid computing outcomes during old-age retirement). A three-year tenure restriction is imposed, meaning that only workers who also had their main place of employment at the closing establishment in t_{-2} and t_{-3} are considered.[30] Finally, we require that the individual is observed in the population register until at least t_1 so as to have information on their outcomes. Those who emigrate or decease before t_1 are therefore excluded.

Our strategy imposes tight restrictions on what constitutes an admissible event. We ignore other adverse events, such as mass layoffs where fewer than 90 percent of an establishment's employees lose their jobs, in order to get as clean an experiment as possible.[31] The final data set includes around 180,000 displaced workers.

We define a set of *control* workers on whom we impose identical restrictions in all dimensions. Thus, we require an establishment size of at least 5, worker age of 24 to 60, and three years of tenure. However, they must be employed at establishments which survive until year t_1 (retaining at least 10 percent of their original size). As in the case of the displaced workers, we do not impose any other restrictions on what these workers do beyond t_{-1}, nor on their future outcomes.[32]

---

[29] Workers who are involved in *false closures* also cannot be included in the control group for year t_0, as we cannot conclusively establish what happens to these workers' establishments.

[30] This is to ensure that the workers are sufficiently connected to the closing establishment, and is in line with literature conventions (a three-year cutoff is used by e.g. Davis and von Wachter, 2011).

[31] Note that we include the few workers who remain within the original establishment in the cases when the establishment did not fully disappear, but where employment did decline by more than 90 percent. Also, we do not place any restrictions on what the workers do during year t_0. This reduces potential endogenous selection due to early leavers from declining establishments.

[32] Some studies have conditioned the control sample on never being displaced in the post-period. This has been shown to result in overestimation of displacement losses (Krolikowski, 2018).

2.2 Worker, Industry and Location Characteristics

We use various worker, location and industry characteristics in our analysis. Here, we provide an overview of the data with details in Appendix A.

The characteristics are grouped into different blocks. Data on basic *demographics* include age, gender and indicators for being a first or second generation immigrant. *Family status* includes indicators for married/cohabiting and divorced, the individual's share of total household labor earnings, number of children in total and number of children of school age. The latter captures any additional impediments to mobility that may arise when children start school at age 7. This block of variables also includes measures of the strength of ties to the worker's current location in the form of a dummy for being born outside of the current county, and the number of times the worker has moved across local labor market boundaries in the past 10 years.

General human capital is captured by years of schooling, labor market experience (years employed during the last ten years), pre-displacement earnings rank among the population of displaced and control workers, and earnings in the years t_{-3} and t_{-2} relative to the year t_{-1}. *Specific human capital* is intended to capture the degree to which the worker is tied to the closing firm or sector, which may determine how costly or difficult it is for the workers to switch industry. We use information on firm tenure and tenure in the same industry as the closing firm (both truncated at 10) and education specificity.[33] We also include indicators for STEM education and for licensed occupations (e.g. nurses).

Characteristics of the *lost job* are important, because the impact of displacement may depend on lost match quality and rents. The data we use include plant size in the year before displacement, the trend in plant size, a manager dummy, and the wage premium associated with the closing plant.[34] We also account for the routine task component of the lost job, by exploiting occupational data and routine intensity measures based on the Dictionary of Occupational Titles (used by e.g. Autor and Dorn, 2013; Goos et al., 2014).[35] Finally, we generate an industry-education match indicator for workers who were employed in one of the 10 main industries of their educational field, and a measure of the size of the displacement event as a share of total employment in its industry-location cell.[36]

[33] We characterize the specificity of education using data on the 1-digit level and 3-digit field of the highest achieved education. Our specificity measure uses the fraction of workers, by level-field combination, that is employed in the top ten 3-digit industries for that field.

[34] The wage premium at the closing establishment was shown to be particularly important in the case of Austria by Gulyas and Pytka (2019). We measure it as the residual wage conditional on demographics, education and industry.

[35] Routineness was found to be an important predictor of post-displacement losses by Blien et al. (2021) and Yakymovych (2022).

[36] The latter is motivated by previous studies, e.g. Cederlöf (2019) and Gathmann et al. (2020), that have found that the impact of being displaced in a large event may be particularly severe.

[14] Measured as the employment growth in the industry from t_{-1} to t_3.

Two blocks of variables relate to aggregate conditions in the worker's industry and local labor market. There is evidence that workers have comparative advantages in their industry of employment and that shocks to this industry have an impact on their overall earnings prospects (Carlsson et al., 2016; Lamadon et al., 2019). Displacement may also have more long-lasting negative effects in industries with low labor turnover (e.g. manufacturing) than in fluid sectors. Our *industry* characteristics (at the 3-digit level) include measures of the average churning rate using conventions from Burgess et al. (2000), and the excess reallocation as the excess creation and destruction of jobs over what was needed to adjust employment in the industry. We also measure the long run industry employment trend over the past 10 years, the current industry-specific business cycle conditions,[14] as well as indicators for the manufacturing sector and the education, health and public administration sectors. These public sector industries have a constantly high demand for workers, which makes it likely that displaced workers are less affected (see Eliason, 2014b). Finally, we include a measure of the education-adjusted wage premium in the industry.

As with industries, there is ample evidence suggesting that local labor market conditions have causal effects on workers' outcomes (see, e.g., Carlsson et al., 2019). We measure these at the level of *local labor markets*, which are constructed by Statistics Sweden based on commuting patterns. We include the local unemployment rate and the population density. Local exposure to industry-related factors is also considered. This is done by constructing shift-share instruments for local industry trends over the past 10 and coming three years, local industry churning and local industry reallocation rates. Concentration of employment in the local labor market is measured by the HHI index across 3-digit industries. We also include the local share of manufacturing employment.

Finally, we include the year when the worker was displaced and the national unemployment rate as variables. This is because earlier studies have suggested that workers displaced during recessions suffer more than those displaced during booms (Eliason and Storrie, 2006; Davis and von Wachter, 2011). Our array of characteristics gives a uniquely detailed characterization of the worker's individual and aggregate situation when the job was lost.

2.3 Outcomes

The outcome that is the main focus of our study is annual earnings. Earnings are normalised by the amount the worker earned in t_{-1}; this captures the relative effects of displacement for each individual without dropping workers with zero earnings in the post-period. We also consider a binary outcome for whether workers are employed (i.e., whether they earn more than three times the monthly minimum wage during the year). This measure is equal to one in

t_{-3} through t_{-1} for all included workers by construction. Geographical and industry mobility is measured by whether the worker lives in a different local labor market or works in a different industry compared to t_{-1}. Industry mobility can only be measured conditional on the worker being employed.

Outcomes are measured from t_{-3} until t_{10}. Our sampling ensures that all displaced and controls are observed in the data between t_{-3} and t_1. From t_2 on, the outcomes are missing if the individual is not in the population register, because they have become older than 64, moved abroad, or deceased.

2.4 Who Are the Displaced?

As can be seen in Table 6 in Appendix B, about 180,000 workers who meet our criteria were displaced due to the closure of some 22,000 establishments in 1997-2014. Over 4,000,000 workers at more than 200,000 establishments are eligible as controls over the same period.

There are significant differences in the composition of the displaced and control groups. In particular, the displaced are concentrated in manufacturing industries; fewer workers are displaced from industries such as education, health and public administration. This industry mix is connected to the overrepresentation of men among the displaced. Another characteristic of displaced workers is the smaller size of their plants, which is in line with well-known facts about the higher survival rates of larger plants.

2.5 Pre-Matching

Because of the appreciable differences in the characteristics of displaced and control workers illustrated in Table 6 in Appendix B, we employ a matching procedure to ensure covariate balance. This is done by propensity score matching on the full set of covariates included in the analysis. We use a logit function to estimate propensity scores and drop workers who lie outside of the common support region.[37] We then match three control workers to each displaced worker without replacement.

2.6 Estimated Average Effects

As a first step of our empirical analysis, we consider the average effects of displacement on normalized annual earnings, employment, and location and industry mobility using our matched sample. The results are estimated as mean differences in outcomes between the displaced and controls and are shown in Figure 1.

[37] Common support in terms of covariates among the displaced and controls is an assumption of GRF estimation.

The short-run impact on earnings and employment is large, and the effects are strongly persistent. Displaced workers are far from regaining the levels of control workers even ten years after displacement. As can be seen in Figure 1a, earnings of displaced workers drop some 24 percent relative to controls by t_1, and are still 8 percent lower in t_{10}. Figure 1b shows that the very persistent earnings effects are mainly driven by a significant drop in employment probability.

The impact on geographical mobility (living in a different local labor market) is modest at around 1.5 percentage points, but clearly significant (Figure 1c). The difference between displaced and controls is entirely due to events during the first two post-displacement years. This suggests that job loss induces a group of workers *who would otherwise not have moved* to relocate shortly after displacement. This is evidence that the impact is not just an adjustment in the timing of mobility.

The results for industry mobility, measured conditional on being employed, in Figure 1d show a somewhat different pattern. Workers become much more likely to switch industry when displaced, which is unsurprising as they need to find a new employer. Half of the workers find a new employer in another 3-digit industry than the one they were displaced from, with 30 percent even switching 1-digit industries. Over time, this impact gradually declines, as more of the control workers switch industries. However, even 10 years after displacement, the difference in the probability of working in another 3-digit industry remains at almost 30 percent.

Figure 1: *Average effects of displacement on earnings, employment status, and geographical and industry mobility*

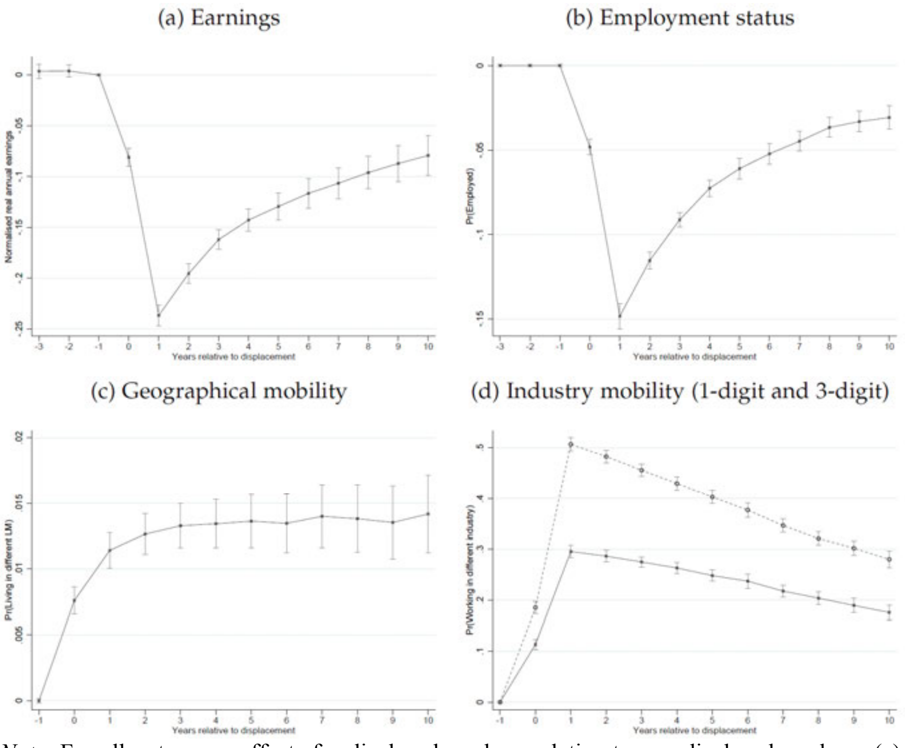

Note: For all outcomes, effects for displaced workers relative to non-displaced workers. (a): Annual earnings (normalized by earnings in the year before displacement). (b): Employment status defined as having annual earnings higher than three times the 10th percentile-level monthly wage. (c): Geographical mobility is an indicator for living in a different local labor market than in t_{-1}. (d): Industry mobility is an indicator for working in a different 1-digit (solid line) and 3-digit (dashed line) industry than in t_{-1}, conditional on being employed in the period. 95% confidence intervals shown.

3. GRF Estimation and Calibration

3.1 Causal Forest Estimation

Our goal is to understand heterogeneity in displacement losses across different groups of workers. We do this using causal forest estimation (Athey et al., 2019), which was designed specifically for understanding heterogeneous treatment effects. Causal forests are a machine learning method which relies on the results of a large number of recursive causal tree algorithms, developed by Athey and Imbens (2016). Each of the causal trees splits the sample of workers into two groups based in turn on each possible threshold level of each included worker characteristic. For every way in which it splits the sample, the algorithm computes treatment effects within the two resulting worker groups. The split that yields the largest difference in estimated treatment effects is selected. The two resulting groups of workers are recursively split again according to the same criterion. This sorts the workers, based on their characteristics, into "leaves" with similar estimated treatment effects.

While a single causal tree finds the splits which best capture treatment effect heterogeneity in the sample considered, the estimates of single trees can be nonrobust. A causal forest resolves this issue by making use of the output of many trees. Each tree is estimated on a random subsample of the workers, and only a random subset of the characteristics is considered when evaluating each split. Because of this, trees are not identical to one another, and the causal forest reveals relationships that hold consistently across random subsamples of workers. The causal forest estimates treatment effects for each worker based on the outcomes of neighbours who end up in the same leaf in the forest's trees. We refer to these predicted treatment effects as conditional average treatment effects, CATEs. In our setting, a large CATE means that the predicted earnings loss for a worker is small (i.e., less negative).

A key advantage of causal forests relative to traditional sample splitting approaches is the data-driven way in which characteristics and thresholds are selected. This leaves little room for the potentially arbitrary choices that researchers might otherwise make when choosing which groups to compare. Another important strength of the forest is its mitigation of overfitting, i.e., the problem of finding spurious heterogeneity when testing across many different splits of the data. This is achieved by constructing trees using the approach known as honesty, and by employing 5-fold estimation (both are discussed in more detail below). Finally, the forest is very flexible when it comes to capturing nonlinear effects of worker characteristics and high-order interactions.

3.2 Implementation of GRF

We define treatment status as binary, corresponding to whether the individual was displaced during a plant closure in the year t_0. The outcome used for estimating the causal forest is annual earnings in the year t_1 after displacement, normalised by annual earnings in t_{-1}. We include all covariates listed in Section 2.2 and discussed in more detail in Appendix A.

For efficiency (see Nie and Wager, 2021), the displacement status W_i and size of loss y_i are made orthogonal to the vector of observed worker characteristics x_i before estimating the causal forest. Two separate regression forests[38] are constructed to estimate the conditional propensity score $\hat{e}_i = W_i|x_i$

[38] The regression forests (Breiman, 2001) are similar in spirit to the causal forest. One regression forest predicts the probability of displacement as a function of the covariates x_i, while the other predicts the earnings loss as a function of x_i (without considering treatment status). The data are split into "prediction trees", each of which allocates the sample of workers into "leaves" based on their characteristics so as to best predict their displacement status or t_1 earnings. The

and marginal response function $\widehat{m}_i = y_i | x_i$. The treatment status and the outcome are residualized to obtain $\widetilde{W}_i = W_i - \hat{e}_i$ and $\tilde{y}_i = y_i - \widehat{m}_i$. The causal forest is then estimated using these residualized values.

We employ 5-fold estimation with a held-out test set. The training set consists of all displaced workers from 80 percent of the closing establishments, as well as their matched controls (the matching procedure is described in Section 2.5). The training set is divided into 5 folds, containing an equal number of closing establishments. Each fold is left out in turn, and a causal forest is estimated on the remaining folds. This forest is then used to predict CATEs for the left-out fold. No information about a worker or his establishment is therefore used when predicting that worker's CATE. This minimizes the risk of overfitting, i.e. capturing idiosyncratic patterns in the data. All ranking of workers according to their CATEs is done within each of the five folds separately. The 20 percent of closing establishments which are in the test set together with their workers' matched controls are used for evaluating different targeting policies in section 7. The test set is also divided into five folds; each of the causal forests estimated using the 5-fold procedure in the training set is used to predict CATEs for one test set fold.

All sampling and splitting conducted by the causal and regression forests' trees is clustered at the establishment level. This is important, because that is the level at which the shock of displacement takes place. Sampling workers from the same establishment into different trees and using them to estimate each other's CATEs could lead to overfitting.

The workers drawn when constructing each of the forest's trees are randomly divided into two halves, one of which is used for determining which splits to make, and the other for estimating CATEs within the resulting leaves. This is known as honesty and serves to mitigate overfitting. Honesty ensures the estimates' consistency and asymptotic normality (Athey et al., 2019).

4. Distribution of Heterogeneous Effects

4.1 GRF Output and Calibration

Figure 2 contains a histogram of CATEs for workers in the training set, as estimated by GRF. The outcome is annual earnings in the year t_1 relative to the t_{-1} baseline. The histogram reveals considerable variation in displacement losses, with some worker groups predicted to cope with displacement relatively well. Others are severely affected, with a long tail of workers who suffer large losses extending from the median to the left. The CATEs are negative

regression forest combines the output of many trees, with each tree being estimated on a random subsample of workers. The forest then predicts the probability of displacement or the t_1 earnings for each worker.

for all workers; there is no worker group for whom displacement is estimated to lead to higher earnings.

Figure 2: *Distribution of causal forest predictions of displacement losses in terms of normalised earnings in the year* t_1

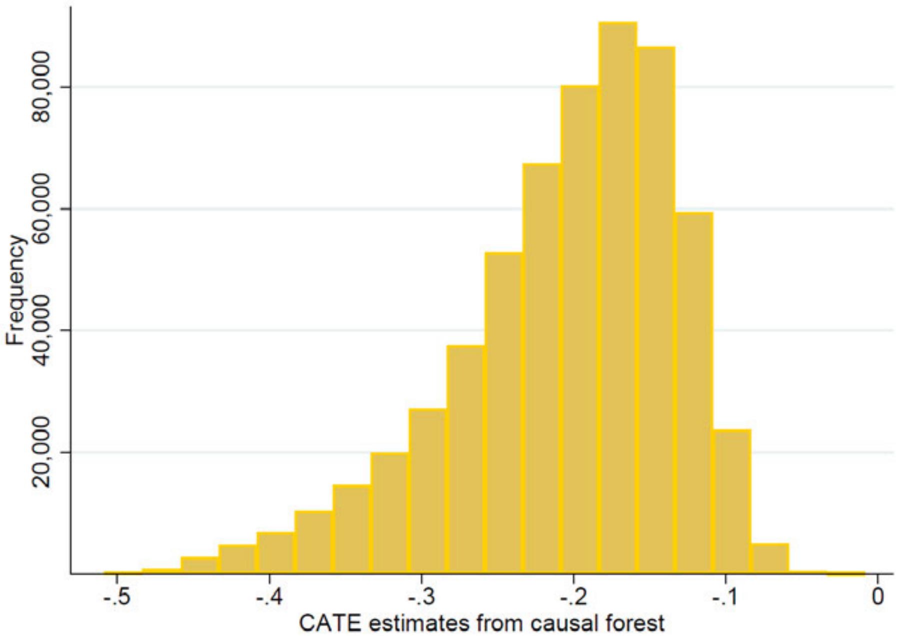

Note: Frequency distribution of predicted causal forest CATEs for the training sample of workers. The outcome is labor earnings in the year after displacement (t_1) as a fraction of pre-displacement earnings in the year t_{-1}.

In most of this paper, we only use these estimated CATEs of displacement to classify workers, industries and locations for further analysis. We then calculate displaced-control differences within these groups and use these as our baseline estimates of displacement losses. We sometimes refer to such within-group displaced-control differences as ATEs, in constrast to the CATEs, which are produced by the GRF. However, it is still necessary to understand whether the CATEs are well-calibrated.

The most formal way of testing the forest's calibration is by using the best linear predictor test (Chernozhukov et al., 2020). It assesses whether GRF can predict both the average treatment effect and variations around this average effect correctly. The causal forest's CATE estimates $\hat{\tau}_i$ and the regression forests' orthogonalised earnings outcome $\tilde{y}_i = y_i - \hat{E}(y_i|x)$ and displacement probability $\widetilde{W}_i = W_i - \hat{E}(W_i|x)$ are used for this purpose. The $\tilde{y}_i$ are regressed on a function of $\widetilde{W}_i$ and $\hat{\tau}_i$ as follows:

$$\tilde{y}_i = \alpha(\bar{\tau}\,\widetilde{W}_i) + \beta\left((\hat{\tau}_i - \bar{\tau})\,\widetilde{W}_i\right) + \varepsilon_i, \qquad \bar{\tau} = \frac{\sum_i \hat{\tau}_i}{N}$$

The parameter α estimates how well the forest's average predicted treatment effect fits the data. If the GRF's prediction of the average displacement effect is correct, then $\alpha = 1$. The parameter β measures if heterogeneity in treatment effects is adequately captured. If β is significantly different from zero, the null of no heterogeneity in treatment effects can be rejected. Optimally, $\beta = 1$. If $\beta < 1$, there is overfitting by the GRF, as its predicted deviations from the mean are larger than the actual deviations from the mean. If $\beta > 1$, the forest does not capture all of the heterogeneity present. This omnibus test results in $\beta = 1.57$ and $\alpha = 1.06$, with both being significantly different from zero at conventional levels. This means that the null hypotheses of no displacement effect and of no heterogeneity in displacement effects are rejected. The value of β implies that the degree of heterogeneity is underestimated. However, this is a much smaller concern than overfitting would be.

4.2 Heterogeneity in Terms of Earnings

To illustrate the heterogeneity that we find, workers are sorted into deciles based on their estimated displacement effects from GRF (CATEs). Those who suffer the most are placed into Decile 1 and those who are the most resilient into Decile 10.[39] We then calculate average effects for displaced workers in each decile by taking the observed difference between the displaced and controls in the decile.[40] We label these average effects ATEs and plot them in Figure 3. It shows that there is substantial heterogeneity in effects of jobb loss, and that the GRF has successfully identified groups of workers who are more and less affected. Displaced workers in the lowest decile lose 46 percent of their earnings compared to controls. The size of the loss decreases monotonically when moving up the deciles, and is only five percent of earnings in Decile 10.

These estimates can also be used to assess the degree of underfitting as identified by the best linear predictor test in Equation 4.1. To this end, the dashed line in Figure 3 plots the average CATEs within each decile for comparison. These underestimate losses in the left tail and overestimate losses in the right tail, with their dispersion more compressed towards the mean than that of the ATEs. This is consistent with the underfitting result of the calibration test. For this reason, we predominantly use the CATEs to rank workers, and rely on ATEs as baseline estimates of displacement losses.

An alternative method of calculating treatment effects within groups of workers is to compute average inverse-probability weighted scores (AIPW).

[39] The division into deciles is done within each of the five folds generated by the 5-fold procedure when implementing GRF estimation in Section 3.2.
[40] We rely on the matching procedure in Section 2.5 to ensure that the displaced and controls in each decile are comparable.

These are based on the displacement propensity and earnings predictions produced by the two regression forests which are run for orthogonalization prior to implementing the causal forest. AIPW has the advantage of giving correct estimates even if only one of these two regression forests is correctly specified. AIPW scores within each CATE decile are also plotted in Figure 3. These are similar to the ATEs; in each case, the AIPW estimate lies within the ATE estimate's 95% confidence interval. The small size of the gains provided by AIPW estimation confirms that our matching procedure has been successful in identifying comparable controls for the displaced workers. For this reason, we calculate ATEs as simple displaced-control differences within worker groups in the remainder of the paper.

Figure 3: *Effects of displacement on normalised earnings in the year t_1, by CATE deciles*

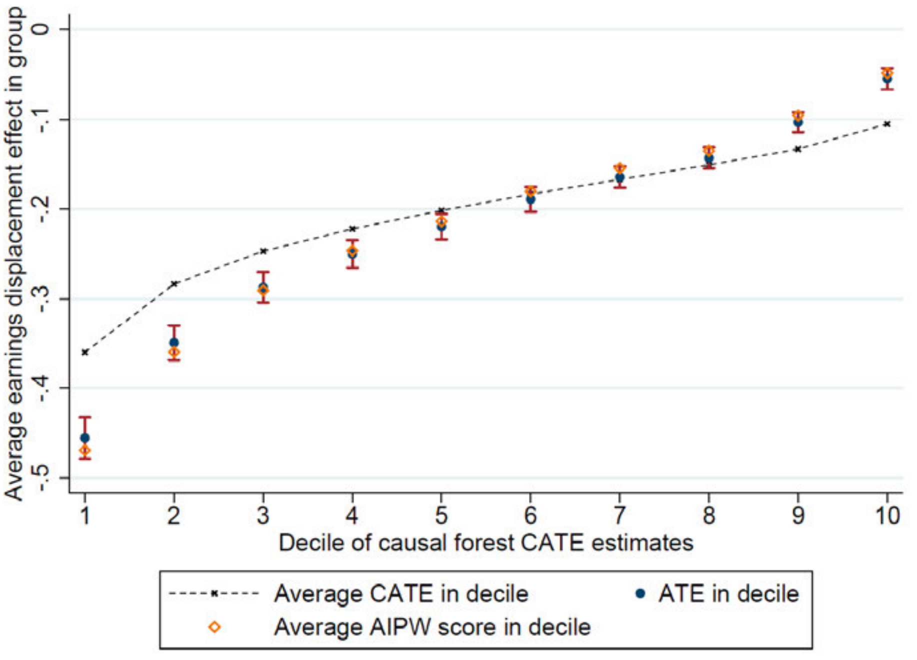

Note: Workers ranked according to treatment effect estimates from GRF and divided into deciles, with Decile 1 containing the most severely affected and Decile 10 the least affected. ATE estimates within deciles estimated as regressions of normalised earnings in the year t_1 on displacement. AIPW scores based on causal forest treatment effect estimates and regression forest estimates of t_1 earnings and displacement propensity. 95 percent confidence intervals of the ATE estimates shown.

4.3 Heterogeneity in Terms of Employment and Long-Run Effects

Even though the causal forest is calibrated to capture heterogeneity in earnings losses in the year t_1, we are also interested in effects on other labour market outcomes. In particular, we consider the binary *employment* measure, which captures whether an individual participates in the labor market during a year.[41] Also, it is important to understand whether heterogeneity in t_1 is transient or whether it can predict the size of losses in later years.

[41] Note that this measure is closely related to the earnings outcome, as it is defined as earning at least three times the minimum monthly wage during the year.

Figure 4: *Effects of displacement on employment in the year t_1, by CATE deciles*

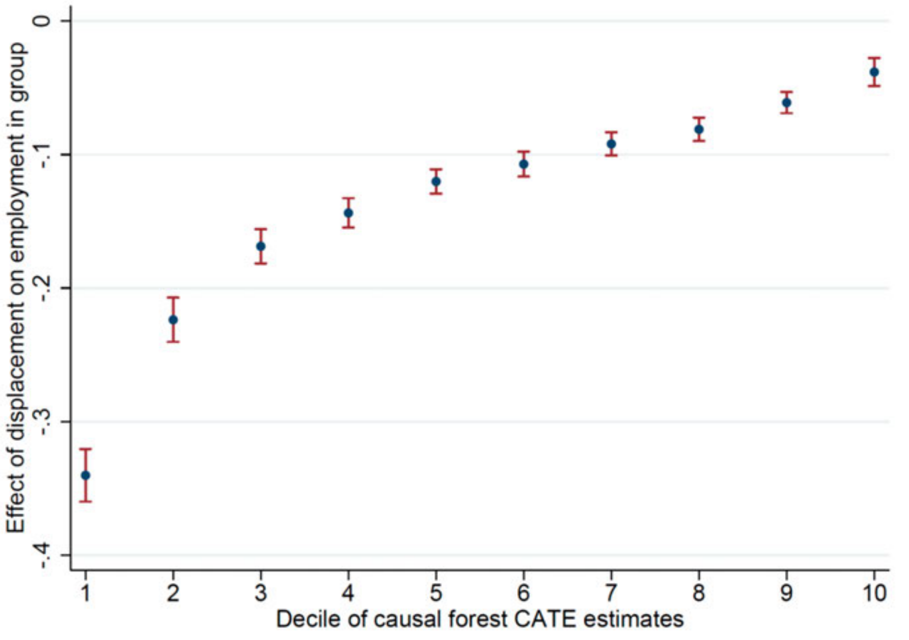

Notes: Workers ranked according to GRF CATE estimates for t_1 earnings and divided into deciles, with Decile 1 containing the most severely affected and Decile 10 the least affected. ATE estimates within deciles estimated from regressions of employment in the year t_1 on displacement. Employment status defined as having annual earnings above three times the minimum monthly wage. 95 percent confidence intervals shown.

We once again rank workers by their GRF CATEs and subdivide them into deciles. The ATEs within each decile are presented in Figure 4. The average displaced worker is 14 percentage points less likely to be employed following displacement, but this masks substantial differences. The GRF is able to capture this heterogeneity remarkably well. Workers in the top decile of CATEs are only four percentage points more likely to be non-employed, compared to 34 percentage points for those in the bottom decile, the effects differing by a factor of 10.

Figure 5: *Earnings and employment effects when splitting the sample using GRF CATEs*

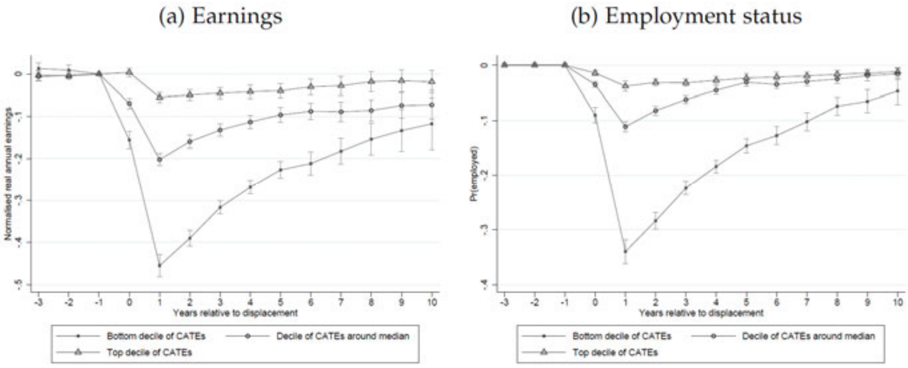

Notes: Effects as the mean difference between the displaced and the corresponding control group workers. The outcomes are annual labor earnings (normalized by earnings in the year before displacement t_{-1}), and employment status, defined as having annual earnings in excess of three times the minimum monthly wage. The decile of CATEs around the median corresponds to those between the 45th and 55th percentiles. 95 percent confidence intervals shown.

Figure 5 presents evidence that the size of displacement effects in the year t_1 is strongly predictive of long-run outcomes. We plot how normalised earnings (Panel a) and employment probability (Panel b) develop from t_{-3} to t_{10} among workers in the top and bottom deciles, as well as among workers in the decile centered on the median CATE.[42] The size of losses in t_1 predicts losses for the entire follow-up period. The earnings losses of the hardest-hit decile in the year t_5 are six times larger than those of the most resilient decile. Employment effects are equally persistent, with the hardest-hit workers having a six times higher probability of being non-employed compared to the least affected ones five years after displacement. There is some convergence in absolute terms towards the end of the period, but the difference between workers in the top and bottom deciles remains statistically significant at the 95 percent level through t_{10}. In relative terms, the earnings differential between the deciles is actually not estimated to decrease at all. Convergence does seems to happen somewhat more quickly when it comes to employment, with the median worker having almost attained the levels of the top decile by t_5. The strong predictive power of t_1 losses on outcomes in following years means that we capture heterogeneity which is relevant for displaced workers' welfare over the medium and long run. Thus, the characteristics analyzed in the sections below do not only have an effect on workers' immediate post-displacement outcomes, but remain relevant over long time horizons.

[42] The sample size decreases after t_4 because we do not have data for years after 2017 and because some workers reach the typical retirement age of 65.

5. Understanding the Heterogeneous Effects of Displacement

5.1 Observed Heterogeneity and Causality

We now shift focus to exploring how worker characteristics and the conditions under which displacement occurs are related to magnitudes of post-displacement losses. The aim is to give a detailed characterisation of workers who are severely affected, as well as of those who prove resilient. As in any exploration of treatment effect heterogeneity, we correlate attributes with estimated effects. The estimated earnings loss within each split of the data should be considered a causal effect of displacement for that worker group, but the splits themselves should not be given a causal interpretation. When comparing displacement effects among workers with, for example, different levels of education, we will i) claim that the effects for each education group should be interpreted as causal and ii) claim that the differences in estimates between the education groups describe differences in causal effects. However, we do not claim that the differences only arise because of education *per se*, as education may be correlated with other important attributes, whereof some may be unobserved.

5.2 Characterizing Workers with Different Magnitudes of Displacement Losses

To understand how different covariates are related to the size of displacement losses, we consider how their average values change for workers with different estimated CATEs from the causal forest.[43] Figures 6 and 7 contain plots of average values of each characteristic against estimated displacement effects. Non-binary variables are standardized by their grand mean and standard deviation, whereas average values of dummy variables are reported at their empirical values. The size of bubbles corresponds to the number of workers who have CATEs of the corresponding size.

[43] To illustrate that the CATEs correspond well to within-group ATEs (i.e. displaced-control differences), we plot CATEs against ATEs in Figure 11 in Appendix B. It is clear that CATEs predict actual displaced-control differences very well, although they are somewhat shrunk towards the average effect of displacement. This follows from the results of the best linear predictor test (Equation 4.1) and from the plots in Figure 3.

Figure 6: *Worker, location and industry characteristics, by estimated CATE*

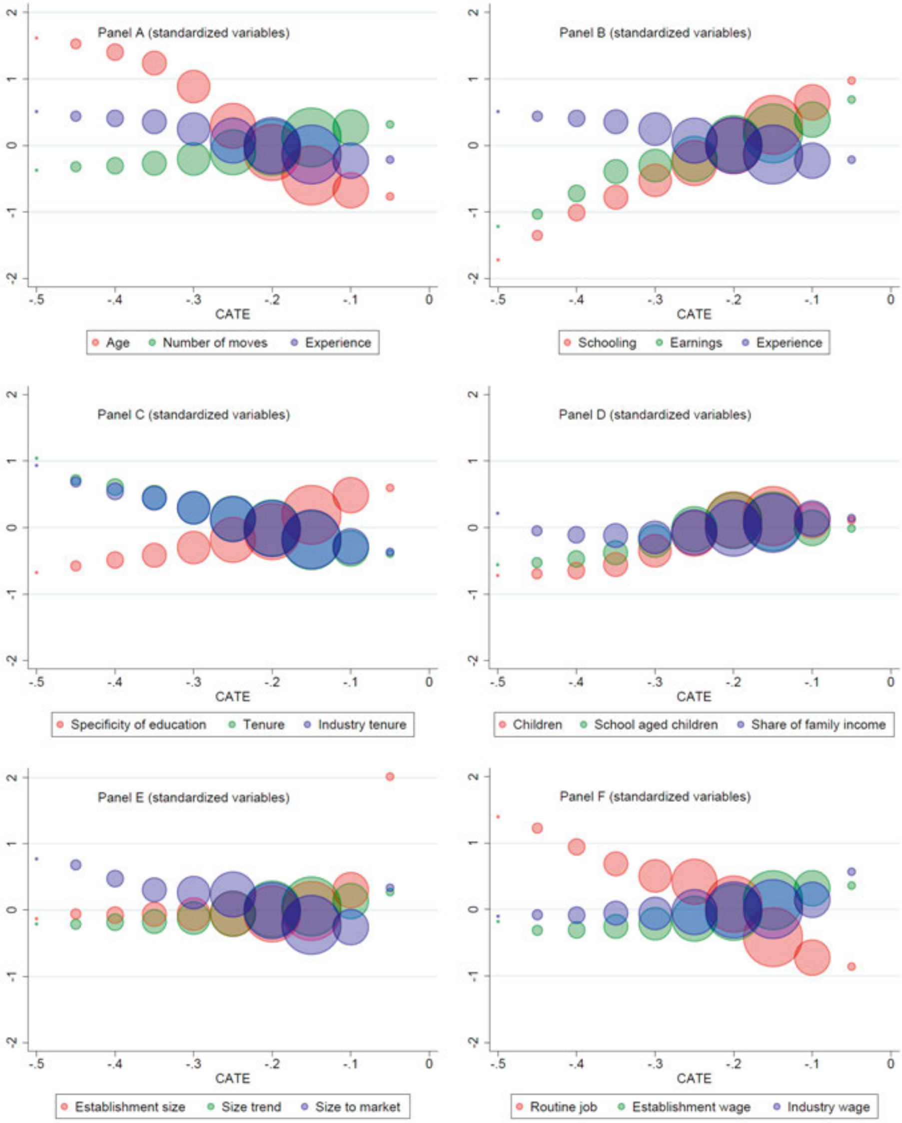

Notes: Average values of variables calculated among groups of workers with different causal forest CATE estimates. Variables standardized by their mean and standard deviation across the full sample of workers. Size of bubbles represents number of workers with given CATE.

Figure 7: *Worker, location and industry characteristics, by estimated CATE (continued)*

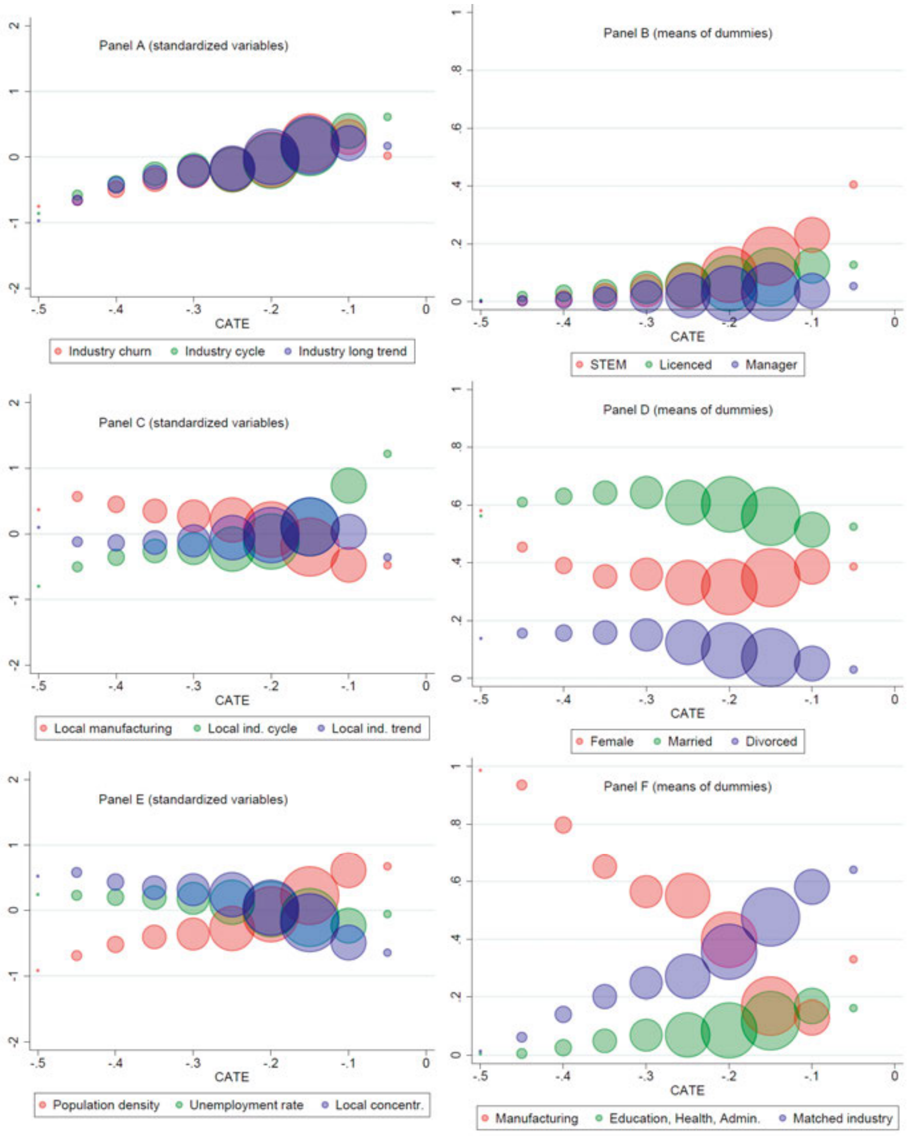

Notes: Average values of variables calculated among groups of workers with different causal forest CATE estimates. Nonbinary variables standardized by their mean and standard deviation across the full sample of workers. Non-standardized values of dummy variables reported. Size of bubbles represents number of workers with given CATE.

Resilient workers tend to be younger and are more likely to have moved across local labor market boundaries in the past. They have more general human capital than non-resilient workers, having more years of schooling and being more likely to hold an education in STEM fields. Resilient workers also tend to have a higher level of earnings before becoming displaced. On the other hand, they seem to have less firm-specific human capital, as evidenced by their shorter job tenures. Non-resilient workers are more likely to have routine jobs. This is related to their concentration in manufacturing industries. They are also exposed to other unfavorable industry characteristics, such as bad long-term and short-term trends, as well as low churn rates. Non-resilient workers are more likely to live in rural areas with high unemployment rates and high manufacturing shares. This translates into their establishment's closure displacing a larger share of workers in their industry-location cell. These results are in line with what is expected based on theory and previous studies.[44]

Plots corresponding to Figures 6 and 7 for individuals' migration background are presented in Figure 13 in Appendix B. Natives who have moved across regional boundaries in the past are expected to do better than those who live in their region of birth. Interestingly, first-generation international migrants are concentrated in the middle of the distribution, being less likely to make it into the top tail, but also somewhat less likely to end up among the most severely impacted workers. Their absence among the hardest-hit workers might be explained by their concentration in large cities.

5.3 Heterogeneity Conditional on Age and Education

Two characteristics which are revealed to be of particular importance in Figure 6 are age and years of schooling. These factors have also been identified to be crucial by previous studies (see e.g., the overview in Kuhn, 2002). Their importance for the causal forest's estimated CATEs is confirmed by Panel A of Figure 8, where average CATEs within age-schooling cells are plotted. The smallest predicted earnings losses are found among workers with post-secondary education in the bottom half of the age distribution. The largest predicted losses are concentrated among the oldest workers without post-secondary education. These predictions are confirmed in Panel B, where displacement effects are estimated as differences between the displaced and controls within each age-education cell. The causal forest's shrinking of the true displacement effect distribution to the mean is evident from the higher dispersion in estimated within-cell ATEs as compared to CATEs (Panel B). The ranking of age-education cells in terms of size of losses is nevertheless the same. Young workers with at least a bachelor's degree lose less than ten percent of their earnings, whereas old workers who have not completed high school suffer losses of over 40 percent.

[44] The importance of these characteristics is confirmed by the variable importance measure presented in Figure 12 in Appendix B. This metric is based on the number of times the causal forest's trees split on a particular variable. Age is the characteristic used for making splits most frequently, followed by plant size, routineness, different location and industry characteristics and education.

Figure 8: *Displacement effects across and within age-education cells*

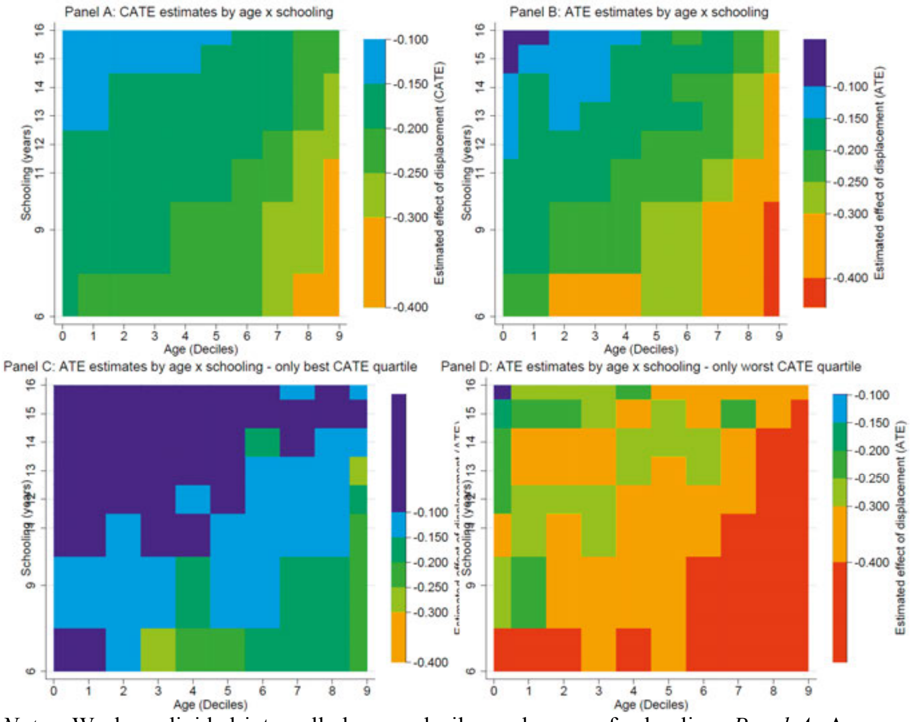

Notes: Workers divided into cells by age deciles and years of schooling. *Panel A*: Average values of causal forest CATEs within each cell. *Panel B*: Regression estimates (ATEs) of differences between displaced and controls within each cell. *Panel C*: Regression estimates of differences between displaced and controls among the most resilient quartile (according to GRF CATEs) within each cell. *Panel D*: Regression estimates of differences between displaced and controls among the least resilient quartile (according to GRF CATEs) within each cell.

In spite of age and education level being important predictors of the size of displacement losses, considerable heterogeneity exists among individuals *within* each age-education cell. This is shown in Panels C and D of Figure 8. Panel C plots ATE estimates among the most resilient quartile within each age-education cell, while Panel D shows corresponding estimates for the least resilient quartile. Differences in the losses experienced by the least-hit and hardest-hit quartiles in each cell are substantial. While the most successful among the young and highly educated do better than the best quartile of the old and low-educated, there are nevertheless groups of low-educated older workers who incur fairly mild losses of under 20 percent of pre-displacement earnings. This is similar to what is experienced by the group of highly-educated young workers as a whole. The worst quartile of relatively well-educated young workers does as badly as older workers with only a compulsory education do on average. The fact that we have been able to correctly identify resilient and non-resilient worker groups *conditional* on age and education is a testimony to the power of GRF estimation.

The existence of substantial heterogeneity within age-education categories raises the question of what factors might underlie it. To investigate possible drivers, we focus on the two extreme age-education groups. These are on the one hand workers under the age of 30 who have attained at least a bachelor's degree and on the other workers older than 50 who have not completed high school. While the first group tends to do better than the second one on average, there is overlap in terms of displacement effects, as can be seen in Panel A of Figure 9. To understand what drives this dispersion, we plot averages of covariates for individuals with different causal forest CATEs *within* these extreme groups in Figures 14-17 in Appendix B.[45] It is clear that industry- and location-specific factors play an important part. Resilient workers are unlikely to work in manufacturing, and are found in industries with relatively high growth and churning rates. They live in densely populated areas, which are home to growing industries, and have low manufacturing shares.

Panels B-D of Figure 9 focus on two prominent characteristics, namely whether a worker is displaced from manufacturing and the population density in the worker's local labor market. In Panel B, the probability of a worker being in manufacturing and the local population density are plotted by GRF CATEs for the old and low-educated group. There are very strong and monotonous relationships between these variables and estimated displacement losses. The most vulnerable workers are exclusively found in manufacturing, while the least exposed ones are almost entirely displaced from non-manufacturing industries. Those with the lowest CATEs tend to live in rural areas, with population densities of about 40; those with high CATEs are concentrated in the dense Stockholm region. Corresponding patterns are found among young highly-educated workers, as evidenced by Panel C. Manufacturing shares drop and population density rises when moving up the CATE distribution. Finally, these patterns can also be seen to hold in the full sample, unconditional on age and education, in Panel D. Interestingly, workers in the very top tail of the CATE distribution in Panels C and D are actually more likely to come from manufacturing industries. This is driven entirely by highly educated workers who have studied STEM fields (i.e. engineers), as illustrated in Figure 18 in the Appendix. This particular group of workers is able to cope fairly well with displacement, in spite of being displaced from manufacturing.[46]

[45] Corresponding plots for migration variables are presented in Panels B and C of Figure 13 in Appendix B.

[46] We would have been unlikely to discover this interesting result using a traditional sample-splitting approach. The causal forest's ability to pick up non-linearities is key to understanding such complex relationships.

Figure 9: *Dispersion of CATEs within extreme age-education groups and their relationship with manufacturing industries and local population density*

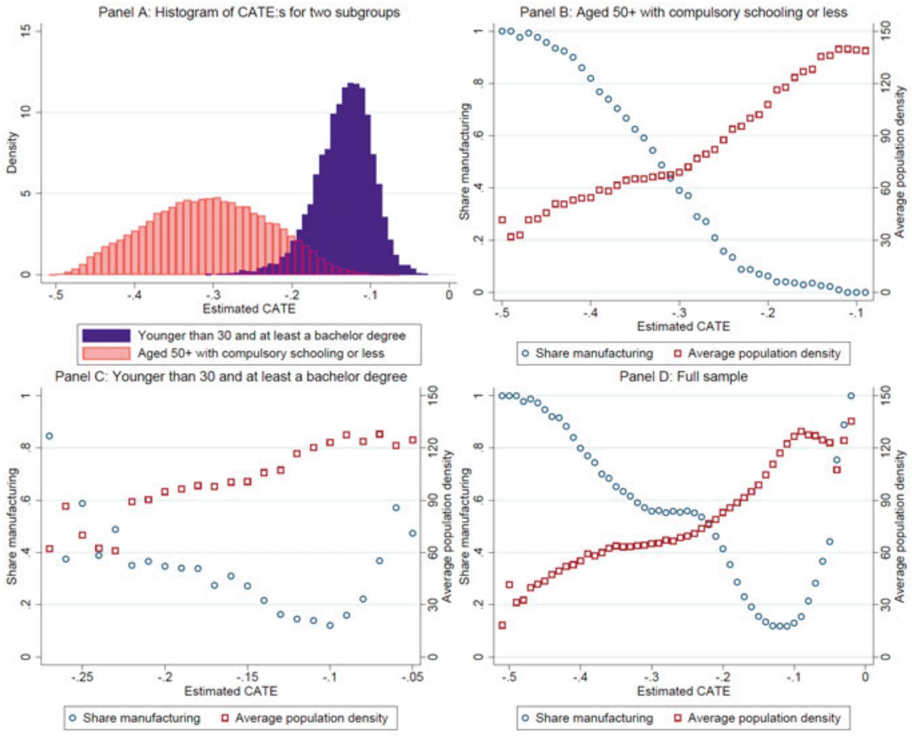

Note: Workers who are (a) younger than 30 and have at least 15 years of education (b) older than 50 and have at most 10 years of education. *Panel A*: Histograms of causal forest CATEs within these groups. *Panels B-D*: Share of manufacturing workers and average population density within each CATE cell, for the two groups and for the full sample. Cells defined as CATE bins of one percentage point. Bins containing fewer than ten workers dropped.

We confirm that the relationships identified in Figure 9 hold by calculating ATEs within groups defined by manufacturing and location density in Table 7 in Appendix B. Among older, less-educated workers, as well as among the general worker sample, manufacturing workers have much higher losses. This is not the case for young, highly-educated workers, as two thirds of manufacturing workers in this group have education in STEM fields. Columns 3 and 4 confirm that STEM-educated manufacturing workers do much better than manufacturing workers without a STEM education. In the full sample of workers, their earnings losses of 19 percent are comparable to the average among non-manufacturing workers, while other manufacturing workers suffer losses of 31 percent. Rural workers lose more than urban ones[47] in all three samples,

[47] We define rural locations as those with a population density of less than 40 persons/km^2, whereas urban locations are those with a population density of more than 100 persons/km^2.

but the gap is more pronounced among older, low-educated workers and less pronounced among young, highly-educated workers.

The strong relationship between location and industry conditions and displacement losses conditional on age and education implies that semi-aggregate factors are important for worker outcomes. We analyze this aspect further in Section 6.

6. Location and Industry Conditions

As evidenced by Figure 9, much of the heterogeneity residual on age and education seems to be related to semi-aggregate conditions at the local labor market or industry level. While population density and manufacturing seem to be the most important predictors of residual displacement losses, it is difficult to consider them in isolation without the other location and industry characteristics. Thus, this section explores combinations of location and industry characteristics as observed to exist in our sample. We classify locations and industries as "good" or "bad" after controlling for differences in worker composition. The results show that workers displaced under bad semi-aggregate conditions suffer much larger loses than those displaced in conditions which are benign. Although rates of location and industry mobility are higher for those displaced in bad conditions, this does not improve their outcomes to match those of workers displaced in good conditions. A possible explanation for this is the very low geographical mobility of all worker groups.

6.1 Classifying Locations and Industries

We characterize how treatment effects vary across different location and industry characteristics separately. A crude way to classify locations or industries as "good" or "bad" would be to rank them according to the average of CATEs within each unit. However, workers are likely to be sorted across regions and industries based on other characteristics. For example, workers who live in an urban region might be younger, more educated and more likely to have an immigrant background than those who live in rural areas. We therefore use the causal forest to estimate treatment effects in each region and industry for *the full population of displaced workers*.

We divide the characteristics X based on whether they capture aspects of the location or industry, or are instead related to the displaced worker or the lost job. The location characteristics are population density, unemployment rate, manufacturing share, short-term and long-term exposure to industry trends and concentration of employment across industries. The industry-level variables consist of dummies for manufacturing and typically public sector industries, churn and reallocation rates, average wages, as well as exposure to forward-looking short-term and backward-looking long-term trends. All other covariates included in the analysis describe features of the individual worker or his workplace at a micro level. The empirically observed vectors of location, industry and micro-level characteristics are denoted X^L, X^S and X^W, with location characteristics in location l denoted X_l^L, industry characteristics in industry s denoted X_s^S, and worker and workplace characteristics of worker w X_w^W. The location and industry where worker w was displaced are denoted $l(w)$ and $s(w)$, with $X_{l(w)}^L$ and $X_{s(w)}^S$ being the corresponding location and industry characteristics.

We estimate partial dependence functions for the location and industry characteristics by setting them in turn to the values observed in each location or industry, while holding the other variables at their empirical levels.[48] We re-calculate displacement effects for these new values using the GRF model and take their mean to evaluate how they affect the size of displacement losses:

$$\tau_l^L = \frac{1}{N}\sum_{w=1}^{N} \text{CATE}(X^L = X_l^L, X_{s(w)}^S, X_w^W),$$
$$\tau_s^S = \frac{1}{N}\sum_{w=1}^{N} \text{CATE}(X_{l(w)}^L, X^S = X_s^S, X_w^W) \qquad (1)$$

τ_l^L and τ_s^S estimate average displacement effects if all workers experienced the conditions in location l or industry s and N is the number workers in the sample. These measures capture the effect of these conditions across the full distribution of workers, making full use of the GRF's non-linear nature. They are not affected by the sorting of individuals across regions or industries, allowing an objective assessment of which of them are associated with larger losses.

6.2 Good and Bad Locations and Industries

We rank locations l by the size of their τ_l^L and industries s by the size of their τ_s^S as estimated by Equation 1. We then split them into quartiles, and focus on the quartiles with the highest and lowest displacement effects. These are referred to as "good" and "bad" locations and industries respectively.

Average values of each location characteristic in good and bad locations are presented in the top panel of Table 1. Good locations are much more densely populated than bad ones, which explains why they hold a larger share of the workers in our sample. They also have lower local unemployment rates, smaller manufacturing shares and a more dynamic and growing industry mix.

[48] This is done within-year, but we suppress the time dimension for presentation reasons. We then average results for the location or industry across years.

In comparison, bad locations are rural areas with high unemployment and industries exposed to structural change. Employment opportunities for displaced workers are likely to be lacking in such areas.

Bad industries are almost exclusively found in manufacturing, while good industries are in non-manufacturing sectors. As expected, this leads to bad industries having lower churning and reallocation rates, as well as declining employment in both the short and the long run.

Table 1: *Sample statistics for good and bad locations and industries*

	Good locations (1)	Bad locations (2)	Diff. [2]-[1]
Population density	109.936	4.842	105.094***
Manufacturing employment share	0.171	0.206	-0.035***
Local unemployment rate	0.080	0.127	-0.047***
Churn rate	0.235	0.202	0.033***
3-year industry employment trend exposure	0.008	0.003	0.005***
10-year industry employment trend exposure	0.086	0.058	0.028***
Reallocation rate	0.150	0.141	0.009***
N workers	436,955	10,060	447,015

	Good industries (1)	Bad industries (2)	Diff. [2]-[1]
Churn rate	0.241	0.170	0.071***
Reallocation rate	0.160	0.106	0.054***
3-year employment trend	0.021	-0.034	0.055***
10-year employment trend	0.243	-0.230	0.473***
Manufacturing	0.006	0.994	-0.988***
Education, health, public administration	0.192	0.000	0.192***
N workers	268,567	86,173	354,740

Note: Location/industry quality measured by placing all workers into all location/industry conditions and predicting CATEs using GRF. Locations/industries are then ranked by average CATE, with good ones being above the 75th percentile and bad ones being below the 25th percentile. Good locations/industries have more workers as they are more populous/larger on average. Significance level: * p<0.10, ** p<0.05, *** p<0.01.

To understand how location and industry conditions affect workers' displacement losses, we first compare the outcomes of workers displaced in good and bad locations and industries. These are shown in Panel A of Table 2. Those who lose their jobs under bad aggregate conditions suffer much more than those who are displaced under better conditions. If displaced in the top quartile of locations, a worker loses 21 percent of their earnings; if displaced in the bottom quartile, a worker loses 1.5 times as much (32 percent). For industries, the difference is even larger, with expected losses in the top quartile amounting to 18 percent, compared to 32 percent in the bottom quartile. However, these gaps might not be due to location or industry conditions, but instead be caused by differences in worker composition across locations and industries.

We would like to compare locations and industries holding the worker composition fixed. To this end, we estimate propensity scores for the probability that a worker ends up in a good or bad location or industry using logit.[49] We then assign individuals weights corresponding to the inverse of these scores and re-estimate displacement effects in Panel B. The estimates do change in a direction which suggests sorting of more resilient workers into better industries (weighted estimates numerically closer than unweighted estimates), whereas the selection of resilient workers into better locations is not very pronounced (weighted and unweighted estimates give similar results). This points to location and industry-specific factors being drivers of displacement losses, rather than just being correlated with adverse individual characteristics.

Table 2: *Displacement effects in different locations and industries*

	All (1)	Good location (2)	Bad location (3)	Diff. [3]-[2] (4)	Good industry (5)	Bad industry (6)	Diff. [6]-[5] (7)
Panel A: Differences between good/bad locations/industries							
Displaced	-0.230***	-0.210***	-0.319***	-0.109***	-0.181***	-0.324***	-0.143***
	(0.005)	(0.005)	(0.022)	(0.022)	(0.004)	(0.028)	(0.028)
Observations	591,324	436,955	10,060	447,015	268,567	86,173	354,740
Panel B: Weighted differences between good/bad locations/industries							
Displaced	-0.230***	-0.212***	-0.314***	-0.102***	-0.188***	-0.297***	-0.110***
	(0.005)	(0.005)	(0.025)	(0.025)	(0.004)	(0.036)	(0.036)
Observations	591,324	436,955	10,060	447,015	268,567	86,173	354,740

Note: The outcome is labor earnings in t_1 (normalized by earnings in t_{-1}). Panel A shows mean earnings differences between displaced workers and matched controls in different location and industry conditions. Panel B weighs the workers in each cell by the probability that someone with their individual characteristics ends up in that cell using IPW. Location/industry quality measured by placing all workers into all location/industry conditions and predicting CATEs using GRF. Locations/industries are then ranked by average CATE, with good ones being above the 75th percentile and bad ones being below the 25th percentile. Good locations/industries have more workers as they are more populous/larger on average. Significance level: * p<0.10, ** p<0.05, *** p<0.01.

6.3 Comparing Importance of Location and Industry Characteristics to Individual Characteristics

To get a sense of how important location and industry conditions are relative to individual characteristics, we *classify workers as resilient and non-resilient based on their individual characteristics*. To this end, we combine the worker characteristics with each combination of the location and industry conditions, and compute for each worker w:

[49] The logit regression is based on demographic, family, mobility and human capital characteristics. It ignores the block of characteristics related to the lost job, as these are closely related to industry.

$$\tau_w^W = \frac{1}{N^L \times N^S} \sum_{l=1}^{N^L} \sum_{s=1}^{N^S} \text{CATE}(X_l^L, X_s^S, X_w^W)$$

where N^L and N^S are the number of locations and industries respectively. The τ_w^W only reflects differences in worker-level characteristics since we place each worker into the same location and industry conditions. We use this measure to classify workers as either resilient or non-resilient.[50]

To illustrate that aggregate conditions matter for displaced workers we compare differences between resilient and non-resilient workers to differences between good and bad aggregate conditions. Columns 1–3 of Table 3 show earnings effects when splitting the sample on the worker dimension. The estimated earnings loss is 14 percent for the most resilient worker quartile and 31 percent for the least resilient worker quartile, a difference of roughly 17 percentage points. Columns 4–6 show that the earnings loss is 17 percent under good aggregate conditions (the top quartile of both locations and industries), and 35 percent under bad aggregate conditions (the bottom quartile of both locations and industries). The difference of 17 percentage points is similar to the one between resilient and non-resilient workers. This supports the notion that the size of displacement losses is determined by both individual characteristics and aggregate conditions.

6.4 Mobility when Faced with Adverse Aggregate Conditions

The fact that conditions at the location and industry level are important for the effects of job loss suggests a clear link to mobility. In the extreme, if workers were perfectly mobile across locations and industries, such aggregate conditions would not matter. On the other hand, it is not clear whether mobility after displacement should matter for the outcomes of workers who are displaced in "good" regions or industries. Displacement does have an effect on geographical and industry mobility, as documented in Section 2.6; here we study whether the impact differs depending on the conditions under which a worker is displaced. The units of analysis are the 72 local labor markets and 211 three-digit industries. Both outcomes are measured three years after displacement so as to capture what happens after workers have had time to adjust to the shock. As we analyze different sets of initial location and industry conditions, all regressions are weighted to capture differences in worker composition according to the procedure described in Section 6.2.

[50] It would be preferable to compute the average of CATEs if a worker were placed in all combinations of location and industry rather than to use the results of Equation 2. However, it is prohibitive in terms of time to re-shuffle each worker into all combinations of location and industry conditions.

Table 3: *Displacement effects across worker and aggregate conditions*

	Resilient worker (1)	Non-resilient worker (2)	Diff. (3)	Good locations Good industry (4)	Bad locations Bad industry (5)	Diff. (6)
Regression	-0.141***	-0.312***	-0.171***	-0.172***	-0.350***	-0.179***
	(0.004)	(0.007)	(0.006)	(0.004)	(0.046)	(0.046)
Observations	295,662	295,662	591,324	207,844	2,249	210,093

Note: The outcome is labor earnings in t_1 (normalized by earnings in t_{-1}). Columns 1–3 split the sample by worker resilience, where resilience is based on each workers' predicted earnings fromm GRF when placed in all workers in all location and industry conditions (see Section 6.3 for details). Columns 4–6 split the sample by location and industry quality. Location/industry quality measured by placing all workers into all location/industry conditions and predicting CATEs using GRF. Locations/industries are then ranked by average CATE, with good ones being above the 75th percentile and bad ones being below the 25th percentile. Good locations/industries have more workers as they are more populous/larger on average. Significance level: * p<0.10, ** p<0.05, *** p<0.01.

Panel A of Table 4 focuses on location mobility. The effect of displacement is larger for workers in bad locations than for those in good locations. This is expected, since workers in bad locations are more likely to have to move in order to find new employment. However, in absolute terms, the size of the effect remains small; geographical mobility is low regardless of where a worker is located. Results for industry mobility, conditional on finding new employment, are presented in Panel B. Those displaced from bad industries are more likely to switch, as expected. Workers in bad locations are more likely to switch industries than workers in good locations. This might be because employment opportunities in these regions are scarce, forcing them to switch industries in order to find new work.

What kind of locations do the displaced move to? Table 5 presents estimates of the effects of displacement on moving to locations with different characteristics. Displacement increases the probability of moving to locations of all types, but it is clear that workers tend to move to locations with better labor market opportunities. The first two columns focus on the population density of the t_3 location compared to the one the worker inhabited in t_{-1}. Displacement increases the probability of moving to a more dense location by 0.75 percentage points, compared to 0.57 percentage points for moving to a less dense location. Similar results are found for the probability of moving to regions with lower and higher local unemployment in columns 3-4. Columns 5-6 consider whether workers move to locations which are estimated to have higher or lower quality (size of τ_l^L) by our GRF model. The estimated effects are in line with those found for population density and local unemployment. Finally, in the last two columns, we investigate whether workers tend to move to areas which are estimated to be high-quality for them personally. This is an aspect which is difficult for individuals to observe, unlike urbanisation and the states of local labor markets. We measure this personalized quality by the size of the displacement loss GRF predicts the worker would have if displaced in the region where he or she moves to compared to the one where he or she was actually displaced. The results do point to workers moving to regions which are a better fit for them personally, but there is some indication that the effect is weaker than for the other, more easily observable regional characteristics. This suggests that workers might not have a full understanding of what conditions favor them personally, but instead consider factors that tend to be good for workers in general.

Table 4: *Location and industry mobility and aggregate conditions when displaced*

*Panel A: Effects on **location** mobility*	All (1)	Good location (2)	Bad location (3)	Diff. [3]-[2] (4)	Good industry (5)	Bad industry (6)	Diff. [6]-[5] (7)
Displaced	0.013***	0.010***	0.045***	0.034***	0.013***	0.017***	0.004
	(0.001)	(0.001)	(0.010)	(0.010)	(0.001)	(0.004)	(0.004)
Control mean:	0.033	0.030	0.061	0.046	0.033	0.033	0.033
Observations	587,309	433,752	10,005	443,757	266,729	85,615	352,344

*Panel B: Effects on **industry** mobility*	All (1)	Good location (2)	Bad location (3)	Diff. [3]-[2] (4)	Good industry (5)	Bad industry (6)	Diff. [6]-[5] (7)
Displaced	0.457***	0.440***	0.561***	0.121***	0.367***	0.533***	0.166***
	(0.008)	(0.008)	(0.023)	(0.024)	(0.006)	(0.049)	(0.049)
Control mean	0.230	0.235	0.214	0.224	0.221	0.216	0.217
Observations	513,434	380,206	8,580	388,786	231,708	74,067	305,775

Note: Location mobility across 72 local labor markets and industry mobility across 211 industries between t_{-1} and t_3. Industry mobility measured only if the worker is employed three years after displacement. Location/industry quality measured by placing all workers into all location/industry conditions and predicting CATEs using GRF. Locations/industries are then ranked by average CATE, with good ones being above the 75th percentile and bad ones being below the 25th percentile. Good locations/industries have more workers as they are more populous/larger on average. Significance level: * p<0.10, ** p<0.05, *** p<0.01.

7. Policy Targeting

Given the potentially severe impact of displacement on workers, policymakers might wish to intervene to ameliorate the effects. The sizeable heterogeneity we find however suggests that supporting all groups of workers equally would be misguided; indeed, some groups of workers only suffer small losses, making interventions unnecessary. This makes correct targeting of policies crucial so that efforts are directed towards the hardest-hit individuals. Such policies might take different forms, including cash transfers, re-training, and subsidies for employers who hire displaced workers. Transfers directly alleviate the income losses displaced workers experience, while the other interventions aim to improve their re-employment prospects. However, the strong correlation between the impact of displacement on earnings and employment, as well as between short-run and long-run effects, suggests they should be targeted towards the same groups of workers. We therefore focus on identifying rules that policymakers can use to reach individuals with large losses in terms of t_1 earnings.

Table 5: *Displacement and moving to different types of locations*

	Population density		Local unemployment rate		Location quality		Personalized location quality	
	More dense	Less dense	Lower unemployment	Higher unemployment	Better location	Worse location	Better location	Worse location
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Displaced	0.0075***	0.0057***	0.0077***	0.0055***	0.0077***	0.0055***	0.0071***	0.0061***
	(0.0006)	(0.0005)	(0.0007)	(0.0007)	(0.0006)	(0.0006)	(0.0006)	(0.0005)
Observations	587,309	587,309	587,309	587,309	587,309	587,309	587,309	587,309

Note: In Columns 1–2, the outcome is an indicator for moving to a more/less dense location between t_{-1} and t_3. Columns 3–4 compare mobility to locations with lower/higher unemployment rates. Columns 5–6 examine mobility to better/worse locations, when ranking locations by the average predicted CATE after placing all workers into its local conditions. Columns 7–8 defines better/worse based on predicted *individual CATEs for each worker* in the locations where they lived in t_{-1} and t_3. Standard errors in parentheses. Significance level: * p<0.10, ** p<0.05, *** p<0.01.

Firstly, we validate the performance of the GRF model when it comes to identifying hard-hit workers in the held-out test set (which contains workers from 20 percent of the closing establishments, as well as their matched controls). Secondly, we use the test set to explore how close to GRF targeting it is possible to get using simple policy rules. Such rules might be necessary to employ in practice because policymakers are unlikely to have access to all the information on which the GRF model draws, making its highly individualized, nonparametric estimates unfeasible to estimate. Furthermore, it can be unethical to target interventions based on some of the characteristics used in the GRF, such as immigrant background or gender.

Figure 10 assesses to what extent different policy rules succeed in reaching hard-hit individuals. We consider situations where between five and 25 percent of those in our sample are targeted. The blue line shows the displacement effect among the selected workers if targeting is random. As expected, the earnings loss among this randomly selected group is close to the test set average of 24.6 percent. Selecting workers with the most negative GRF CATEs identifies a group of workers which is much more severely impacted, as shown by the maroon line. The five percent of workers who are hit the hardest according to GRF experience displacement losses of 54 percent. Among the 25 percent hardest-hit, earnings losses are 41 percent. This shows that our model has been able to properly identify test set workers who are strongly impacted by displacement, even though it was constructed without using them in any way. The GRF therefore picks up relationships between worker characteristics and sizes of displacement losses that hold in the general population of Swedish displaced workers.

The importance of age and education emphasised in Section 5.2 suggests comparing GRF-based targeting to targeting older workers or workers with low education. The green line in Panel a of Figure 10 shows the size of displacement losses if workers are targeted based on their age, oldest first. Losses are 42 percent among the oldest five percent of workers, and 34 percent among the oldest quartile. Targeting the least-educated workers, shown by the yellow line, reaches workers who are on average less hard-hit than the oldest workers. The five percent least educated have losses of 37 percent, while those in the least educated quartile have losses of 32 percent. Combining age and education information to target the oldest less-educated workers does not yield large improvements (light green line). Panel b considers targeting based on semi-aggregate local and industry characteristics. If manufacturing workers are selected at random, the losses of the workers reached by the policy are around 36 percent. It is possible to improve targeting by selecting manufacturing workers in routine jobs, as shown by the light green line, but not by much. However, selecting older workers in manufacturing improves targeting significantly. The first five percent of individuals selected according to this rule suffer earnings losses of 48 percent; if a quarter of the individuals in the test sample are selected, their losses amount to 39 percent. Targeting workers in locations with low population density (shown by the red line) does not do as well in terms of identifying the hardest-hit individuals as targeting based on manufacturing.

Figure 10: Average effects of displacement on earnings for workers selected by different targeting policies

(a) Targeting based on age and education

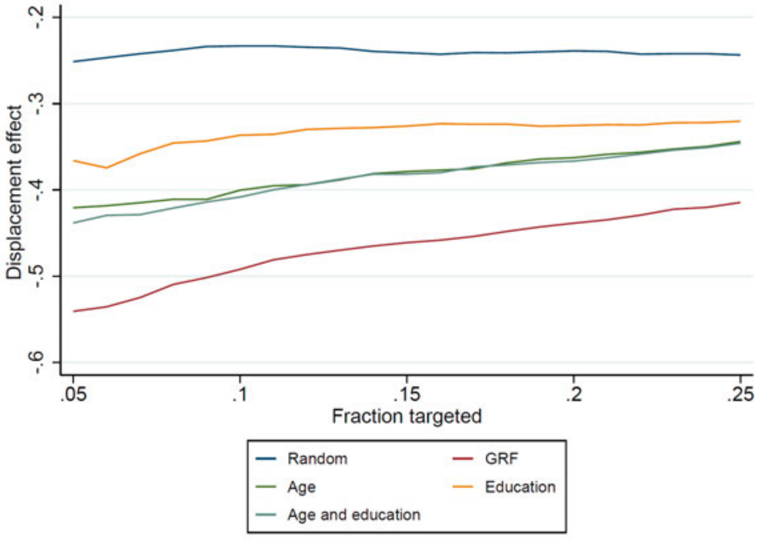

(b) Targeting based on manufacturing and population density

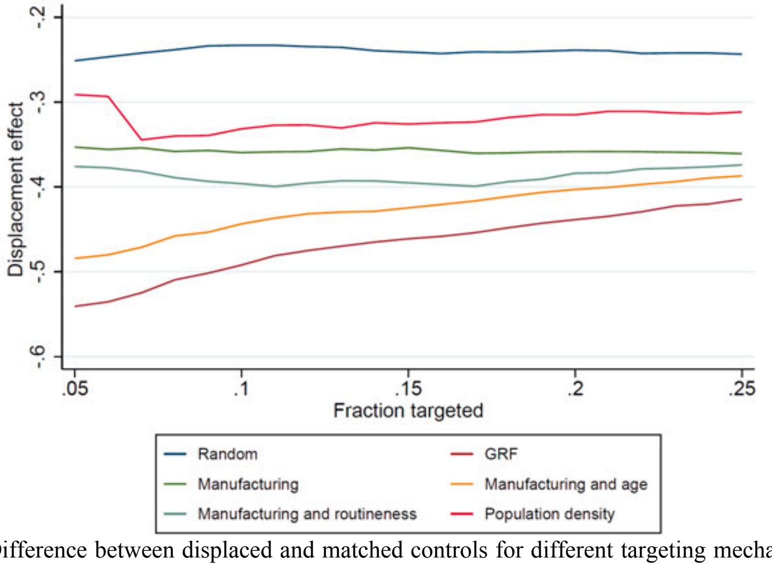

Note: Difference between displaced and matched controls for different targeting mechanisms and targeting shares. Outcome is normalized earnings in t_1. Effects estimated on held-out test set (20 percent of the closing establishments and the matched controls of their workers). *Random*: random set of workers; *GRF*: workers with lowest estimated CATEs; *Age*: oldest workers; *Education*: workers with fewest schooling years; *Age and education*: Orders workers by five-year age bins and schooling years; *Manufacturing*: manufacturing workers at random; *Manufacturing and age*: oldest manufacturing workers; *Manufacturing and routineness*: most routine manufacturing workers; *Population density*: least dense locations.

8. Conclusion

Our paper contributes to a vast literature which attempts to understand why job displacement leads to detrimental outcomes for affected workers. Many explanations for post-displacement losses have been put forward. These have highlighted factors such as firm, industry and occupation-specific human capital (Kuhn, 2002; Neffke et al., 2022; Yakymovych, 2022), firm premia (Gulyas and Pytka, 2019), location mobility (Gathmann et al., 2020) and aggregate conditions (Davis and von Wachter, 2011). Our results confirm that determinants of earnings losses after displacement are multidimensional, as substantial heterogeneity remains even after conditioning on the most important factors. While older workers with lower levels of general human capital lose more on average, even young college educated workers can suffer substantial earnings losses if other circumstances are stacked against them. Aggregate demand-side conditions at the local and industry level play an important part in determining the size of losses, with rural and manufacturing workers more exposed. This pattern suggests that there is no easy rule that policymakers can follow to target retraining resources, or predict where the need for financial assistance will arise after negative shocks. The policy rule which comes closest to achieving the targeting recommended by the causal forest is to focus on older workers in manufacturing. This is a group which suffers from a combination of detrimental individual and semi-aggregate conditions.

References

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

Altonji, J. G., E. Blom, and C. Meghir (2012): "Heterogeneity in human capital investments: High school curriculum, college major, and careers," *NBER working paper*, 17985.

Athey, S. and G. Imbens (2016): "Recursive partitioning for heterogeneous causal effects," *Proceedings of the National Academy of Sciences*, 113(27), 7353–7360.

Athey, S., J. Tibshirani, and S. Wager (2019): "Generalized random forests," *The Annals of Statistics*, 47(2), 1148–1178.

Autor, D. H. and D. Dorn (2013): "The growth of low-skill service jobs and the polarization of the US labor market," *American Economic Review*, 103, 1553–97.

Bartelsman, E. J. and M. Doms (2000): "Understanding Productivity: Lessons from Longitudinal Microdata," *Journal of Economic Literature*, 38, 569–594.

Black, S. E., P. J. Devereux, and K. G. Salvanes (2015): "Losing Heart? The Effect of Job Displacement on Health," *ILR Review*, 68.

Blien, U., W. Dauth, and D. H. Roth (2021): "Occupational routine intensity and the costs of job loss: evidence from mass layoffs," *Labour Economics*, 68.

Bonhomme, S., T. Lamadon, and E. Manresa (2019): "A Distributional Framework for Matched Employer-Employee Data," *Econometrica*, 81(3), 699–739.

Breiman, L. (2001): "Random forests," Machine learning, 45(1), 5–32.

Burgess, S., J. Lane, and D. Stevens (2000): "Job Flows, Worker Flows, and Churning," *Journal of Labor Economics*, 18(3), 473–502.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018): "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 36, S13–S70.

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

Carlsson, M., I. Häkkinen Skans, and O. N Skans (2019): "Wage Flexibility in a Unionized Economy with Stable Wage Dispersion," *IZA Discussion Paper*.

Carlsson, M., J. Messina, and O. N. Skans (2016): "Wage adjustment and productivity shocks," *The Economic Journal*, 126, 1739–1773.

Cederlöf, J. (2019): "Saved by Seniority? Effects of Displacement for Workers at the margin of Layoff," *Job market paper*.

Chernozhukov, V., M. Demirer, E. Duflo, and I. Fernandez-Val (2020): "Generic Machine Learning Inference on Heterogeneous Treatment Effects in Randomized Experiments, with an Application to Immunization in India," *National Bureau of Economic Research, Working Paper*, No. w24678.

Davis, S. J. and T. von Wachter (2011): "Recessions and the costs of job loss," *Brookings Papers on Economic Activity*.

Eliason, M. (2012): "Lost jobs, broken marriages," *Journal of Population Economics*, 25, 1365–1397.

Eliason, M. (2014a): "Alcohol-related morbidity and mortality following involuntary job loss: Evidence from Swedish register data," *Journal of Studies on Alcohol and Drugs*, 75, 35–46.

Eliason, M. (2014b): "Assistant and auxiliary nurses in crisis times: Earnings, employment, and income effects of female job loss in the Swedish public sector," *International Journal of Manpower*, 35, 1159–1184.

Eliason, M. and D. Storrie (2006): “Lasting or latent scars? Swedish evidence on the long-term effects of job displacement,” *Journal of Labor Economics*, 24, 831–856.
Gathmann, C., I. Helm, and U. Schönberg (2020): “Spillover effects of mass layoffs,” *Journal of the European Economic Association*, 18, 427–468.
Goos, M., A. Manning, and A. Salomons (2014): “Explaining job polarization: Routine-biased technological change and offshoring,” *American Economic Review*, 104, 2509–26.
Gulyas, A. and K. Pytka (2019): “Understanding the Sources of Earnings Losses After Job Displacement: A Machine-Learning Approach,” *CRC TR 224 Discussion Paper Series*.
Huttunen, K. and J. Kellokumpu (2016): “The effect of job displacement on couples’ fertility decisions,” *Journal of Labor Economics*, 34, 403–442.
Jacobson, L. S., R. LaLonde, and D. Sullivan (1993): “Earnings Losses of Displaced Workers,” *American Economic Review*, 83, 685–709.
Krolikowski, P. (2018): “Choosing a control group for displaced workers,” *ILR Review*, 71, 1232–1254.
Kuhn, P. J., ed. (2002): *Losing Work, Moving On: International Perspectives on Worker Displacement*, W.E. Upjohn Institute for Employment Research.
Lachowska, M., A. Mas, and S. Woodbury (2020): “Sources of Displaced Workers’ Long-Term Earnings Losses,” *American Economic Review*, 110.
Lamadon, T., M. Mogstad, and B. Setzler (2019): “Imperfect Competition, Compensating Differentials and Rent Sharing in the U.S. Labor Market,” *Working Paper, University of Chicago*.
Leighton, M. and J. D. Speer (2020): “Labor market returns to college major specificity,” *European Economic Review*, 128.
Neffke, F., L. Nedelkoska, and S. Wiederhold (2022): “Skill Mismatch and the Costs of Job Displacement,” *CESifo Working Paper*.
Nie, X. and S. Wager (2021): “Quasi-oracle estimation of heterogeneous treatment effects,” *Biometrika*, 108(2), 299–319.
OECD (2019): “Society at a glance: OECD Social Indicators”
Yakymovych, Y. (2022): “Consequences of job loss for routine workers,” *IFAU working paper*.

Appendix A: Data Details

Below, we explain the different blocks of variables used in the analysis.

Demographics

Data on basic demographics are drawn from administrative population registers. *Age* is measured in years in November of each year. *Gender* is coded by a female dummy. A categorical *immigrant* variable takes the value zero for natives, 1 for second generation immigrants and 2 for first generation immigrants. It is derived from information on own (and parents') country of birth. First-generation immigrants are those who are born outside of Sweden unless both parents are born in Sweden (thus, adopted are treated as native-born). Second generation immigrants are non-immigrants both of whose parents are born outside of Sweden.

Family Status and Social Ties to Location

Civil status is coded through two dummies. *Married* workers are within a formal marriage, or cohabit with a partner with whom they have a common child. This definition is well in line with Swedish perceptions of "marriages" and is used in most research on Swedish data. *Divorced* is a dummy for workers whose marriage has ended and who are not currently defined as married according to *Married*.

To quantify the economic value of the partnership and degree of dependence, our *income share* variable measures the subject's labor earnings as a share of total household labor earnings.

We measure the number of children in the household by two different variables. *All children* contains the number of all children under the age of 18 in the household, whereas *school children* instead counts children aged 7 to 17. We make this distinction to allow the second variable to measure the additional impediments to mobility that may arise when children start school at age 7.

Social ties to the current location are measured by two variables. First, we define a *dummy for being born outside of the current county*. This variable naturally takes the value 1 for all foreign-born.[51] We further measure the *number of residence moves* across local labor market boundaries during the past 10 years. This variable is a lagged counterpart of our key indicator for geographical mobility after the event.

General Human Capital

Years of schooling are based on the highest achieved level of education. *Labor market experience* is measured as the number of years employed during the 10 years prior to displacement. Because of the tenure restriction, all displaced

[51] The register information is based on the county of birth and cannot be disaggregated further.

workers must have been employed for at least three years and the variable therefore has a range from 3 to 10.

In this category, we also include pre-displacement earnings. We measure these by *earnings rank* among the full population of displaced and controls in year t_{-1}, as well as the change in earnings from t_{-3} and t_{-2} to t_{-1}. The rank form of the first variable has been chosen to ensure orthogonality to time trends.

Specific Human Capital

It is likely that workers differ in how costly or difficult it is for them to switch firm and industry if hit by a negative shock. The ability to adjust along this margin is potentially determined by the degree to which the worker is tied to the closing firm or sector. In order to measure the empirical relevance of these aspects, we include variables capturing pre-closure establishment *Tenure* (truncated at 10) and the number of years spent in the same industry as the closing establishment (*industry tenure*, also truncated at 10).

Furthermore, we characterize the *specificity of education* using data on the 1digit level and 3-digit field of the highest achieved education (thus, high school programs and college majors). Our metric uses the fraction of workers, by education cell, that is employed in the ten main (3-digit) industries. The strategy follows Altonji et al. (2012).[52] Examples of fields with high levels of specificity are pharmacists and nurses.

We further include an indicator for *STEM* education beyond high school. We also provide a dummy for types of education that cater for the health or education sector. The reason is that many (although not all) of these professions are formally licensed. Thus, we refer to the dummy as *Licensed*.

The Lost Job

The causal impact of displacement for displaced workers will not only depend on their outside opportunities, but also on all aspects of the job that they just lost. We therefore define a number of variables capturing key aspects of the lost job.

We first characterize the lost job by *plant size* in terms of employment in t_{-1}, and *plant size trend* between t_{-3} and t_{-1} measured as the difference in log(*plant size*).

In addition, we measure the *wage premium associated with the closing plant*. This feature was shown to be particularly important in the case of Austria studied by Gulyas and Pytka (2019). They studied mass layoffs and characterized the affected firms by the firm wage effects as estimated conditional on person-fixed effects in the tradition of the AKM model (Abowd et al., 1999). Indeed, there is a large literature discussing the origins of firm-specific

[52] An alternative is provided by Leighton and Speer (2020) but their approach requires a full matrix where all types of education are present in all industries and is therefore less well suited to our granular data.

rents, and how they should be estimated and interpreted (see, e.g., Card et al., 2018; Bonhomme et al., 2019; Lachowska et al., 2020, in the context of plant closures). A drawback of the AKM model is that it requires structural assumptions that we do not want to impose, and is likely to result in biased premium estimates for dying firms. We therefore take a slightly different route. As a reduced form measure of the pay level of the closing firm at the time of displacement, we use the leave-one-out mean of residual earnings in year t_{-1} as in Card et al. (2013) and followers. For each displaced worker and potential control, we characterize the closing establishment environment by the average Mincer residual earnings of all co-workers at the same establishment. Residual earnings are computed through year-specific regressions of log earnings on 3-digit industry indicators, years of schooling, field of education (2 digits), gender, immigration status and a full set of age dummies.[53]

We measure the *routine task component* of the lost job. To this end, we use data on routine intensity by occupation from Autor and Dorn (2013) and Goos et al. (2014). We translate the occupational codes into the Swedish nomenclature. We have data on occupations for about half of our workers. We impute routine intensity for each worker based on the average routine score for those in the same detailed education-industry combination. [54]

We further construct a *manager* dummy for workers who are employed as managers. This information is drawn from the occupational codes, which we do not have for all workers, and therefore contains false negatives. To mitigate the problem, we impute manager status based on data from the previous three years if the information is missing in t_{-1}.

We calculate the *relative size of the displacement event* as the share of total employment by industry and location combination. This is motivated by previous studies, e.g. Cederlof˙ (2019) and Gathmann et al. (2020), that have found that the impact of being displaced in a large event is particularly severe. One possible explanation for this finding is that workers from the same event may compete with each other for the same job openings, and this type of competition should be particularly problematic if the displacement event is large relative to the industry-specific local labor market.

Finally, we generate a variable for *education-industry match*. This indicator takes on the value 1 for workers who were employed in one of the 10 main industries of their field before displacement (defined in the same way as for education specificity above).

[53] The R^2 in a regression of labor earnings on the plant wage premium measure among the displaced is around 22 percent.

[54] The leave-out correlation between our education-industry-based measure of routineness and routineness as measured directly based on occupation is 0.7. For cells where we have fewer than 100 workers we use data on the 1-digit industry and education (field+level) instead. The reason for not using occupation-level routineness directly is the large number of missing values this would lead to.

Industry

There is ample evidence that workers on average have comparative advantages in their industry of employment and that shocks to this industry have an impact on their overall earnings prospects. Examples include the paper by Carlsson et al. (2016) for Sweden, which shows that that technology shocks have a larger impact on workers' wages if the shocks are shared with other firms in the same industry and that this distinction is entirely driven by workers with fields of education where most job-to-job mobility is within the industry. Similar arguments are made in Lamadon et al. (2019). It is also likely that displacement has more lasting negative effects in industries with low labor turnover such as manufacturing than in fluid sectors such as restaurants. Overall, this suggests that workers who lose their jobs in declining low-turnover sectors will suffer more adverse consequences, in particular if much of their human capital is industry-specific.

In order to quantify the conditions in each industry, we first construct timeconsistent industry indicators at the 3-digit level. This is a non-trivial endeavour as the codes changed three times during our sample period. We start from the SNI2002 codes that our raw data provide for the period 2002 to 2010. Next, we rely on the SNI2002 code reported in 2010 and use it for all later years for those establishments that continue to exist beyond 2010, and conversely use the data from 2002 for establishments that existed prior to that. We refer to these overlapping establishments as stayers. We then use the modal overlap between SNI2002 (as imputed for stayers) and the current codes (SNI69, SNI92 and SNI2007) to fill in SNI2002 codes for non-stayers. This works particularly well for the years and sample we study, as almost all closing establishments existed in 2010 (due to the 3-year tenure requirement) and since most codes remained unchanged between SNI92 and SNI2002.

We measure *industry wage premia* in each 3-digit industry as average earnings in the industry, residual on years of schooling, 2-digit education field, gender, immigrant background and a full set of age dummies. Furthermore, we also include average *churning* and *excess reallocation* rates for each industry. We follow conventions from Burgess et al. (2000), and define establishment-level churning as the number of workers who are hired or separated beyond what was needed for the actual change in employment between two adjacent years. The churning rate is measured as churning relative to the average employment during the two years.[55] Similarly, we calculate excess reallocation in each industry as the excess creation and destruction of jobs over what was needed to adjust industry employment. The reallocation rate is measured relative to the average employment in the two years.[56] We then aggregate

[55] Thus, churning is [(Hires + Separations) - abs(Emp(t) - Emp(t-1))] / [Emp(t)/2 + Emp(t-1)/2] at the establishment-year level.

[56] Thus, reallocation is [(Job Creation + Job Destruction) - abs(Emp(t)-Emp(t-1))] / [Emp(t)/2 + Emp(t1)/2] at the industry-year level.

the establishment-year churning rates and the industry-year reallocation rates to industry-level averages which are constant over time. When computing these numbers, we take care to exclude observations that satisfy the conditions for false closures as discussed above. Note also that the scale of the measures (relative to average employment in the two years) is in approximate percentages, but with a maximum of 2 and a minimum of -2.

We also follow the conventions from this literature when computing the *change in employment over the past 10 years* in each industry. Thus, for displacements in year $t-1$, we calculate $Trend_{Ind,t} = (Emp_{Ind,t-1} - Emp_{Ind,t-10})/(Emp_{Ind,t-1}/2 + Emp_{Ind,t-10}/2)$ which takes the value 2 for newly emerging industries and the value -2 for disappearing industries. This metric bounds cases where some very small industries experience extreme changes (including exits and entries) during the sample period. Similarly, we calculate the *industry-specific business cycle* as the change in employment between year t_{-1} and t_0 using the same metric.

We also add dummy variables for closures in the *manufacturing industry*, as this industry is of specific interest. Furthermore, we add a dummy for closures within *education, health and public administration* as the labor markets in these (mostly public sector) industries tend to experience a constant shortage of workers.

Location

As with industries, there is ample evidence suggesting that local labor market conditions have causal effects on workers' outcomes. This is also the case in Sweden, where wages seem to react to local labor market shocks (Carlsson et al., 2019).

Our local variables are measured at the level of local labor markets, which are an aggregation of municipalities. These are constructed by Statistics Sweden based on commuting patterns.

For each local labor market, we measure the *local unemployment rate* as the number of residents who are registered with the public employment service (a prerequisite for receiving benefits from either the unemployment insurance system or from the municipal welfare system) divided by the size of the local labor force (sum of the registered unemployed and the number of employed as described above). Also, *population density* is measured at the local labor market level. *Concentration of local employment across 3-digit industries* is measured by an HHI index.

In addition, we measure the local exposure to the industry characteristics we discuss above. These shift-share/Bart´ık-style variables are calculated by multiplying each industry's employment share in the local labor market (by year) with the characteristics of that industry and then summing over the industries within the local labor market. We thus identify the *share of employment in manufacturing*, *exposure to long-term trends*, *exposure to business cycles*, *average churning* and *average reallocation*.

It is important to note the fundamental difference between these variables and their industry-level counterparts. Whereas the industry-level data measure characteristics in the industry from which the worker was displaced (and is potentially tied to, if changing industry is costly), the local labor market counterpart measure how exposed the worker is to these attributes if searching at random at the local labor market (which should be more pertinent to workers who are restricted in their mobility).

Appendix B: Additional Figures and Tables

Table 6: *Sample statistics for displaced and control workers*

	Controls (unmatched) (1)	Displaced (unmatched) (2)
N workers	4,351,442	180,753
N establishments	230,625	21,796
Covariates		
Demographics and family		
Age	44.111	42.191
Female	0.478	0.347
Immigrant	0.097	0.129
2nd gen. immigrant	0.028	0.032
Married	0.644	0.588
Divorced	0.101	0.099
Own fraction of family earnings	0.714	0.756
Kids living at home	0.780	0.745
School-aged kids living at home	0.524	0.466
Not living in birth region	0.274	0.250
Location moves in last 10 years	0.164	0.210
General and specific human capital		
Years employed in last 10 years	9.229	8.918
Earnings rank	0.500	0.496
Trend in earnings rank	-0.000	0.004
Years of education	12.162	11.469
Tenure (censored at 10 years)	6.624	5.908
Industry tenure (censored at 10 years)	7.922	7.145
STEM education	0.097	0.109
Licensed occupation education	0.238	0.071
Education specificity	0.565	0.470
Lost job characteristics		
Plant size	521.033	184.207
Plant size trend	0.025	-0.035
Manager	0.033	0.027
Job routineness	0.515	0.558

Plant wage level	0.001	-0.035
Education-industry match	0.560	0.374
Share of establishment in local industry employment	0.170	0.160
Industry characteristics		
Churn rate	0.223	0.209
Reallocation rate	0.131	0.139
3-year trend	0.005	-0.001
10-year trend	0.051	0.078
Manufacturing	0.229	0.371
Education, health, public administration	0.407	0.096
Location characteristics		
Population density	82.468	85.995
Manufacturing employment share	0.189	0.185
Local unemployment rate	0.086	0.087
Churn rate	0.224	0.226
3-year industry trend exposure	0.007	0.007
10-year industry trend exposure	0.080	0.078
Reallocation rate	0.147	0.148
HHI by industries	0.031	0.031
Year	2005.881	2005.574
National unemployment rate	7.765	7.802

Note: Mean values of all covariates included in the analysis for the unmatched samples of displaced and controls.

Figure 11: *Comparison of CATE estimates from the causal forest and ATEs (displaced-control differences) within each CATE bin*

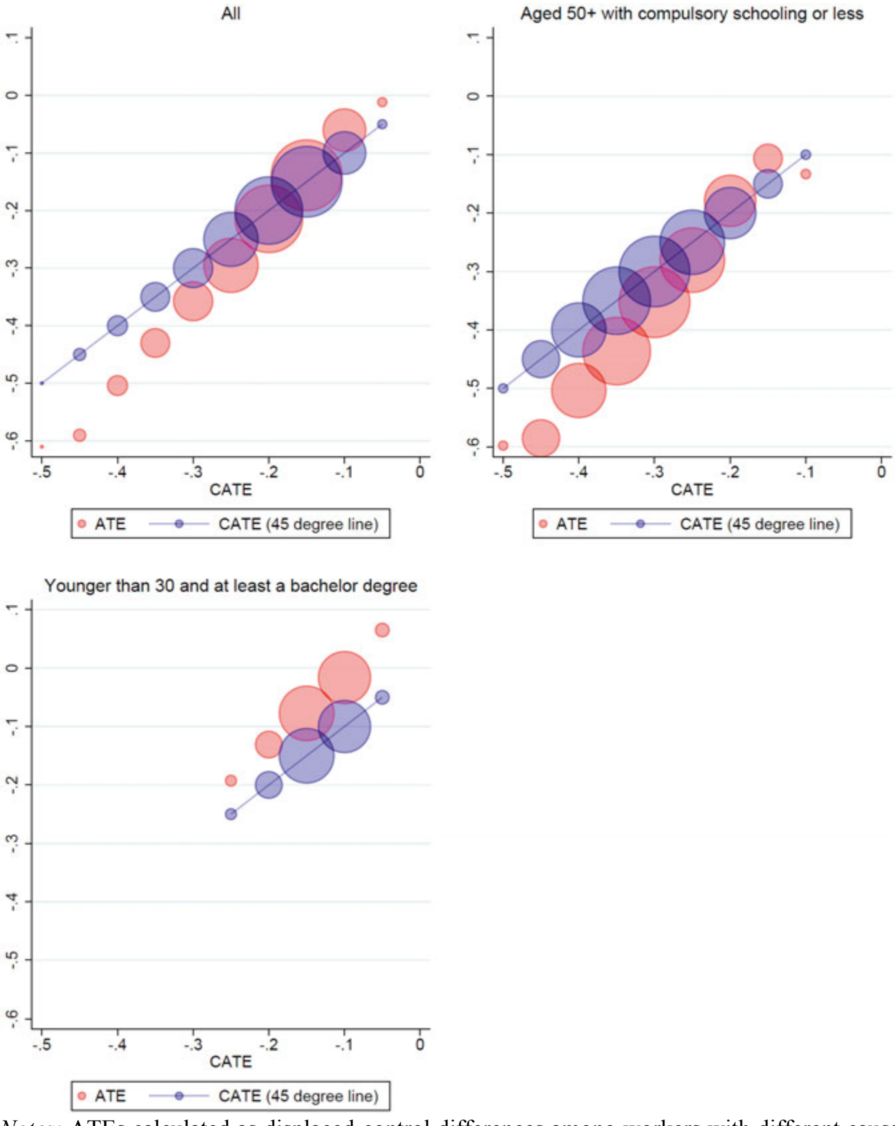

Notes: ATEs calculated as displaced-control differences among workers with different causal forest CATE estimates (sorted into five percentage point bins). Size of bubbles represents number of workers with given CATE.

Figure 12: *Variable importance as measured by fraction of splits made on each variable in causal forest trees*

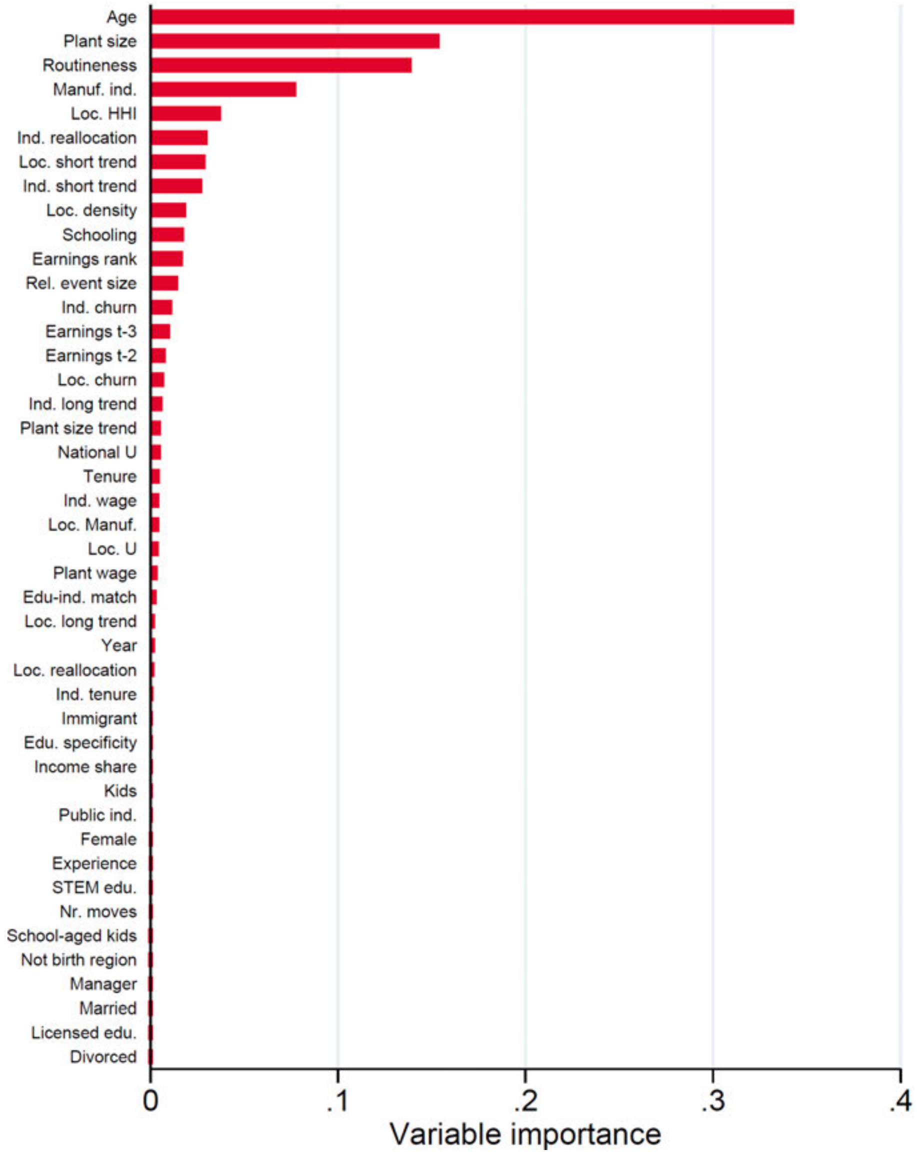

Note: Importance is measured as share of splits at maximum depth of 4 within the trees. Splits at lower depth d given two times the weight of those at $d + 1$. Total importance sums to 1.

Figure 13: *Mean values of migration-related characteristics*

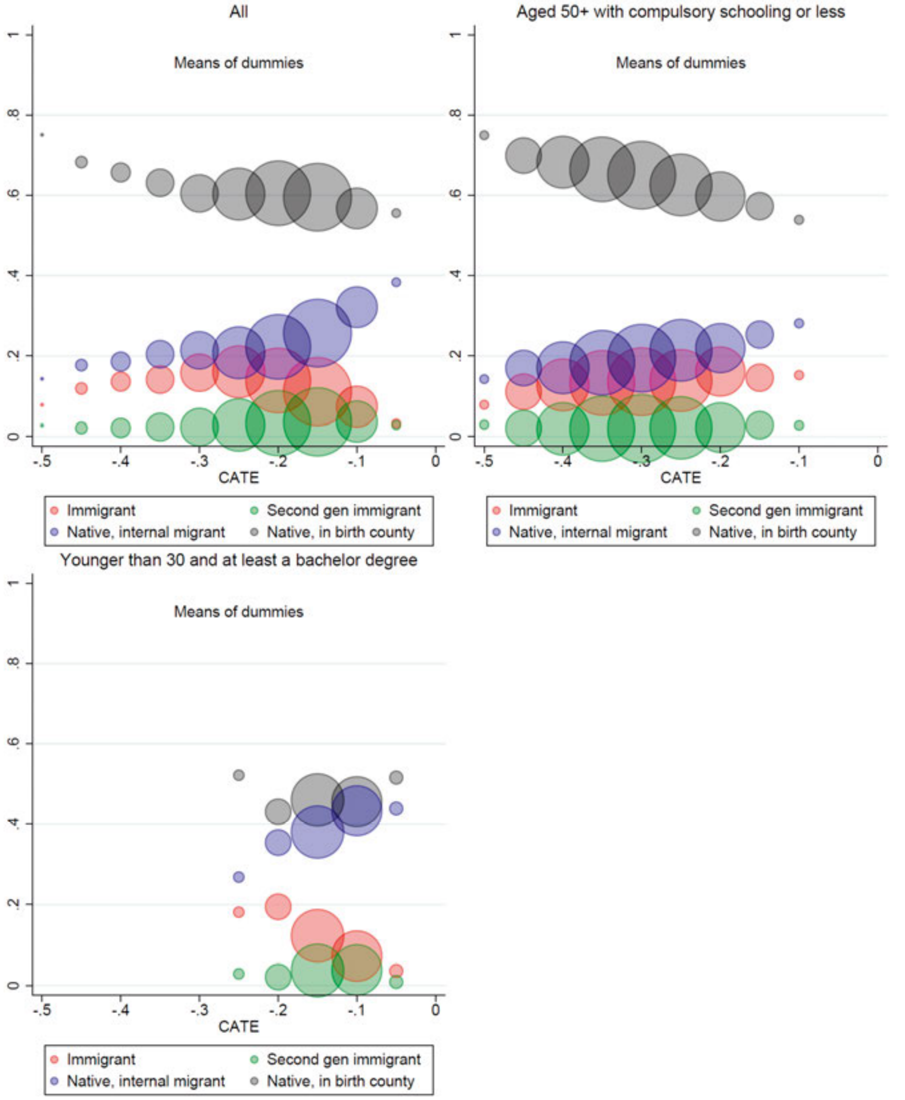

Notes: Average values of variables calculated among workers with different causal forest CATE estimates for the full population, as well as for the old and low-educated and young and high-educated groups. Variables standardized by their mean and standard deviation across the full sample of workers. Size of bubbles represents number of workers with given CATE.

Figure 14: *Mean values of worker, industry and location characteristics for young, highly-educated workers*

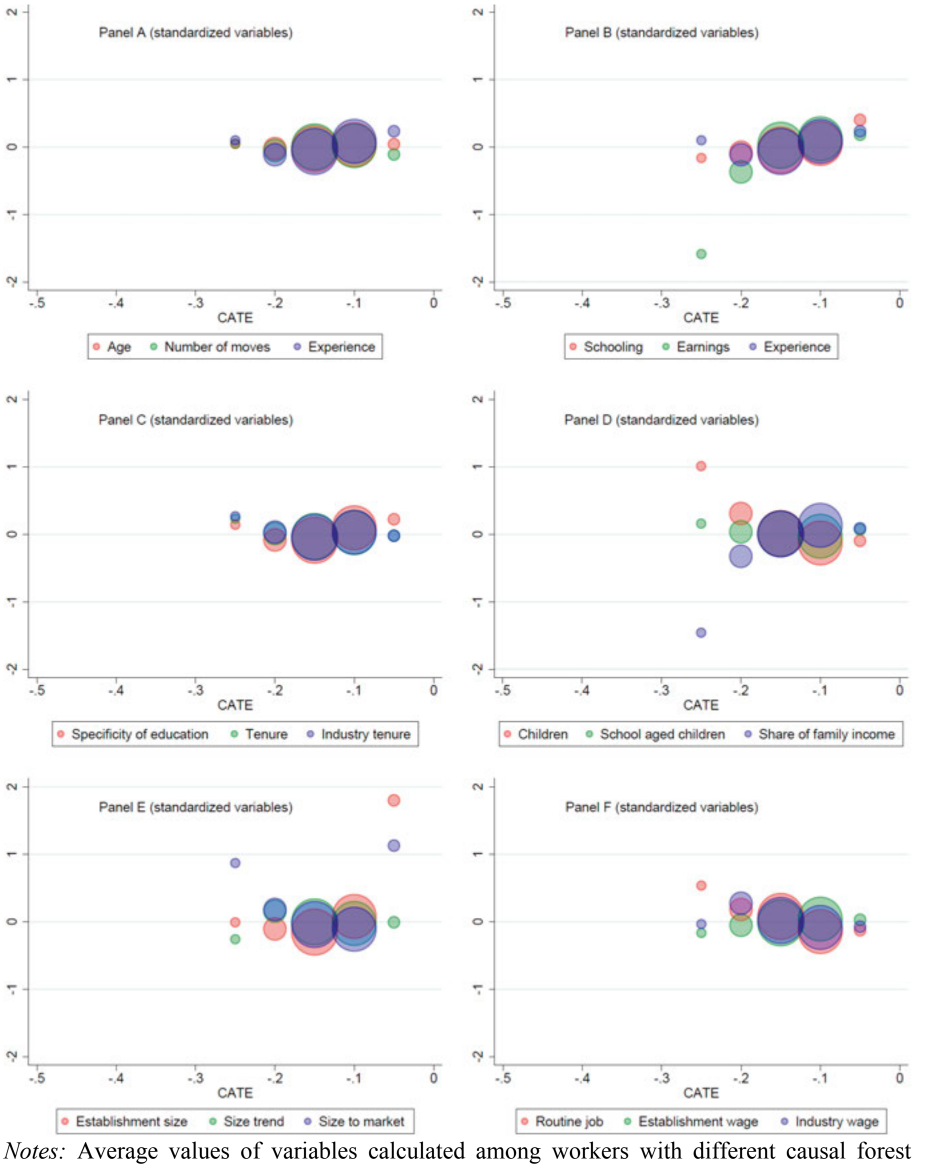

Notes: Average values of variables calculated among workers with different causal forest CATE estimates within the young and highly-educated group. Variables standardized by their mean and standard deviation across the full sample of workers. Size of bubbles represents number of workers with given CATE.

Figure 15: *Mean values of worker, industry and location characteristics for young, highly-educated workers (continued)*

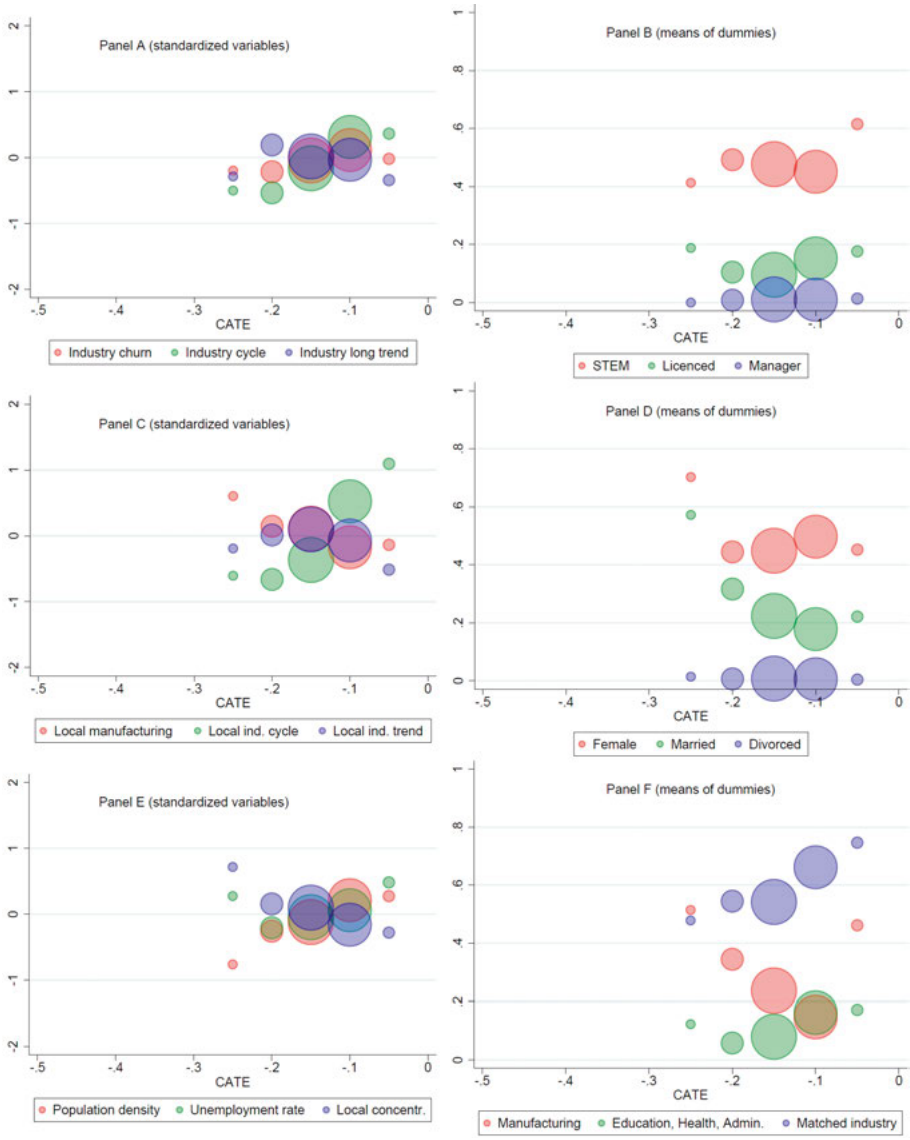

Notes: Average values of variables calculated among workers with different causal forest CATE estimates within the young and highly-educated group. Non-binary variables standardized by their mean and standard deviation across the full sample of workers. Dummy variables not standardized. Size of bubbles represents number of workers with given CATE.

Figure 16: *Mean values of worker, industry and location characteristics for old, low-educated workers*

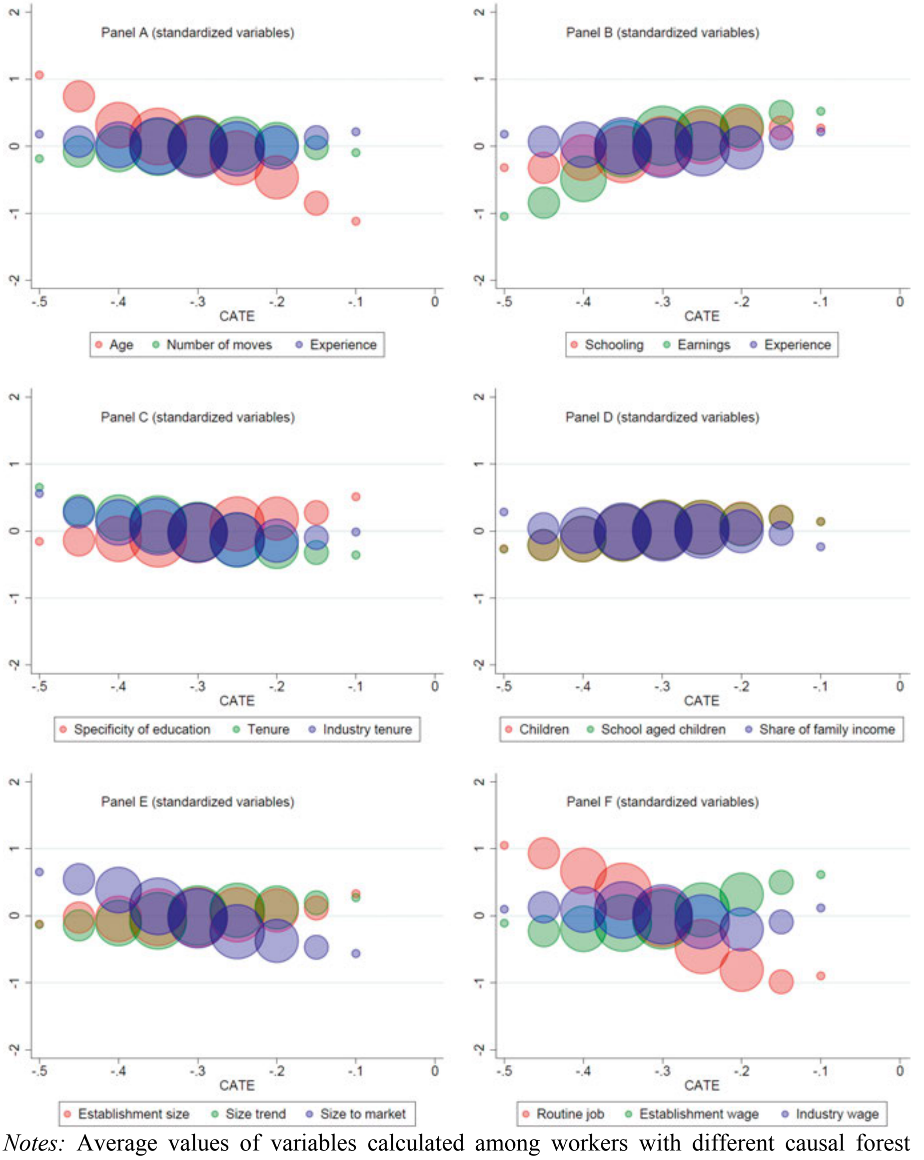

Notes: Average values of variables calculated among workers with different causal forest CATE estimates within the old and low-educated group. Variables standardized by their mean and standard deviation across the full sample of workers. Size of bubbles represents number of workers with given CATE.

Figure 17: *Mean values of worker, industry and location characteristics for old, low-educated workers (continued)*

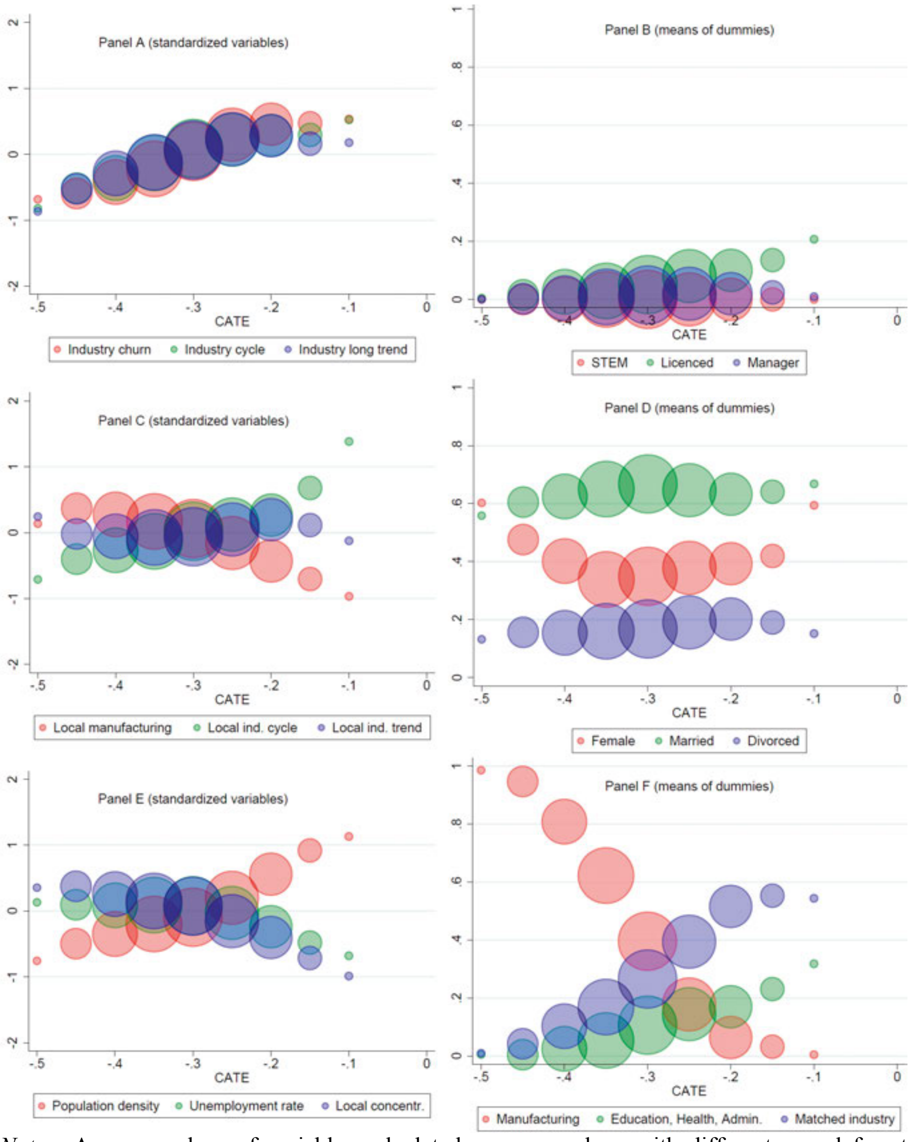

Notes: Average values of variables calculated among workers with different causal forest CATE estimates within the old and low-educated group. Non-binary variables standardized by their mean and standard deviation across the full sample of workers. Dummy variables not standardized. Size of bubbles represents number of workers with given CATE.

Figure 18: *Share with STEM education across CATE bins, in total and in manufacturing*

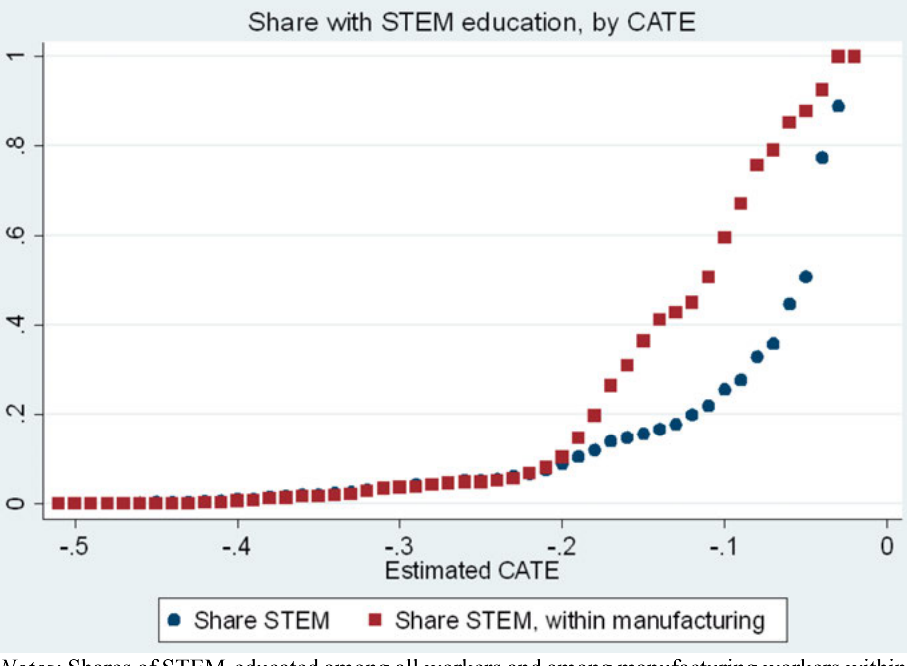

Notes: Shares of STEM-educated among all workers and among manufacturing workers within each CATE cell. Cell size equal to one percentage point of CATEs.

Table 7: *Treatment-control differences among different subgroups of old, low-educated and young, highly-educated workers*

	Non-manuf. (1)	Manuf. all (2)	Manuf. STEM (3)	Manuf. not STEM (4)	Urban (5)	Rural (6)
Panel A: *Old, low-educated workers*						
Displaced	-0.302***	-0.459***			-0.312***	-0.422***
	(0.00766)	(0.0175)			(0.0119)	(0.0127)
Observations	31,099	24,363			22,171	30,796
Panel B: *Young, highly-educated workers*						
Displaced	-0.0639***	-0.0568	-0.0409	-0.0912	-0.0558**	-0.0735**
	(0.0213)	(0.0347)	(0.0387)	(0.0601)	(0.0223)	(0.0339)
Observations	6,365	1,822	1,198	624	5,537	2,442
Panel C: *All workers*						
Displaced	-0.190***	-0.297***	-0.186***	-0.313***	-0.192***	-0.269***
	(0.00350)	(0.0154)	(0.0263)	(0.0132)	(0.00616)	(0.00884)
Observations	372,923	218,401	26,220	192,181	288,863	278,723

Note: Workers who are (A) older than 50 and have at most 10 years of education (B) younger than 30 and have at least 15 years of education. STEM education defined only at post-secondary level. Rural locations have population density <40 persons/km^2, urban locations have population density >100 persons/km^2.

Essay IV. Understanding Occupational Wage Growth

with Adrian Adermon, Simon Ek and Georg Graetz

We thank Michael Böhm, Stefan Eriksson, Peter Fredriksson, Lena Hensvik, Lisa Laun, Oskar Nordström Skans, as well as seminar participants at the Uppsala Center for Labor Studies and the Future of Labor workshop in Berlin for helpful suggestions. Adermon: Institute for Evaluation of Labor Market and Education Policy (IFAU). Ek, Graetz, and Yakymovych: Department of Economics, Uppsala University.

1. Introduction

The past four decades have seen systematic shifts in occupational employment across industrialized countries, with high- and low-paying occupations gaining at the expense of the middle. This is commonly interpreted as reflecting labor demand shifts induced by technological change, consumer demand, or offshoring. However, the impact of such occupation-level demand shifts on the wage structure is far from clear. First, occupations appear to play a minor role in driving changes in wage inequality, at least in terms of descriptive decomposition exercises. Second, occupational employment and wage growth typically do not feature a strong positive correlation. Finally, wage inequality trends differ substantially across countries, while occupational employment shifts are highly similar.[1]

In this paper, we shed light on these puzzles by studying occupational wage growth in Sweden in 1996–2013. Swedish employment shifts are similar to those elsewhere (Adermon and Gustavsson, 2015), but the wage structure is dramatically compressed compared to most other industrialized countries, and growth in inequality has been moderate and episodic (Graetz, 2020). We show that, as elsewhere, occupations do not appear to play an important role in basic decompositions of changes in wage inequality.

However, as has long been recognized, any analysis of occupational demand shifts and wages must address selection problems arising from workers' systematic sorting into occupations (see for instance Roy, 1951; Acemoglu and Autor, 2011; Böhm, 2020). For example, a positive demand shock to computer programmers may manifest itself as an increase in the price paid for a unit of programming output. At the same time, this increased *occupational wage premium* draws in workers from other occupations, who may be less productive than incumbents, thus leaving *observed* wages approximately unchanged.

Our starting point for overcoming the selection problem is to focus on occupation stayers, whose wage growth comes closer to the growth in premia as time-invariant skills are differenced out and the composition of workers is left unchanged (Cortes, 2016). We address attenuation bias stemming from selection on idiosyncratic shocks using the method developed by Böhm et al. (forthcoming).

[1] The polarization of occupational employment in the US and elsewhere has been documented by Wright and Dwyer (2003); Goos and Manning (2007); Autor et al. (2006); and Goos et al. (2014). See in particular Adermon and Gustavsson (2015) for the Swedish case. Goos et al. (2014) provide evidence in favor of a technological explanation. Barany and Siegel (2018) emphasize structural change and consumer demand instead. In a decomposition exercise, Hoffmann et al. (2020) find only a minor role for occupations in driving rising wage inequality. Roys and Taber (2019), Böhm (2020), and Böhm et al. (forthcoming) highlight the lack of a strong positive correlation between occupational employment and wage growth in the US and Germany.

The second challenge we face is that occupations may differ in how workers accumulate skills over the life-cycle, so that differential wage growth among occupation stayers may reflect not only differential premium growth (Deming, 2021). Moreover, occupational experience profiles may have shifted over time, for instance due to technology-induced obsolescence of skills (Deming and Noray, 2020). A theoretically motivated restriction that has been suggested as a solution to this identification problem is the concept of a "flat spot", a point in the life-cycle when the derivative of human capital with respect to experience is zero (Heckman et al., 1998; Bowlus and Robinson, 2012). We propose a novel approach for implementing this restriction, namely to re-center the experience profiles around the flat spot. This leaves us with greater statistical power as we are not forced to restrict the sample to workers near the flat spot. More importantly, it allows us to estimate experience profiles for each occupation and point in time.[2]

Finally, we explore to what extent our estimated premium growth is driven by changes in occupation-specific skill returns. A growing literature documents changing skill returns in the aggregate, and suggests that occupations may be important drivers of such trends (Deming, 2017; Edin et al., 2022). Given the availability of cognitive and psycho-social skill measures from the Swedish military enlistment, we are able to control for differential changes in skill returns in our estimation.

Our findings are as follows. First, premium growth is positively correlated with employment growth (and more strongly so than is raw wage growth). Second, premium growth is also positively correlated with initial wages. These two findings together imply our third finding, namely that in the absence of compositional changes between-occupation wage inequality would have increased more than it actually has. Fourth, experience profiles vary strongly across occupations at any given point in time, and while they are stable in some occupations, in others they show large changes. These results are robust to allowing for changes in specific returns to cognitive and psycho-social skills.

The positive association between premium growth and employment shifts suggests that variation in premium growth is mostly due to demand side factors. At the same time, our results suggest that there is an important life-cycle component to shifts in the occupational wage structure.

Our findings are consistent with a recent and growing literature documenting the importance of compositional changes in counteracting occupation-level demand shifts (Cortes, 2016; Böhm, 2020; Cavaglia and Etheridge, 2020; Böhm et al., forthcoming). Our contribution compared to these studies

[2] Böhm et al. (forthcoming) assume that experience profiles are constant – following much of the theoretical literature on task-biased technological change – and use a pre-period of uniform premium growth to estimate these profiles. Our data do not go back in time sufficiently to make this approach feasible.

is first, to provide comparable evidence for the Swedish economy, which at first glance features a very different wage structure. Second, to estimate time-varying occupation-specific experience profiles. And third, to allow for time-varying occupation-level skill returns when estimating wage premium growth.

To the best of our knowledge, the joint estimation of premium growth and experience profiles has only been attempted by one other paper, Böhm et al. (forthcoming). Their identification assumption is that the profiles are fixed over time, and that during the decades prior to 1985 any differential wage premium growth was negligible, so that experience profiles can be estimated using a prior period. Our assumptions and identification strategy differ from Böhm et al. (forthcoming), and we view our approach as complementary.[3]

The remainder of the paper is structured as follows. Section 2 presents our theoretical framework, discusses identification challenges as well as our proposed solutions, and develops counterfactual scenarios. We describe our data in Section 3. Section 4 contains our results, and Section 5 concludes.

2. Theoretical Framework and Empirical Strategy

The theoretical motivation for our empirical exercise is the standard Roy model in which workers sort into occupations based on comparative advantage. Rather than estimating a completely specified model, our point of departure is an assumption about the data-generating process for potential wages. In Section 2.1, we explore how key parameters of this wage equation can be identified under different assumptions about occupational choice. In Section 2.2, we show how changes in overall wage inequality can be attributed to occupation-level driving forces, and develop counterfactual scenarios based on our estimated wage equation.

2.1 Identifying the Parameters of the Wage Function

Suppose that individual worker i's log wage in occupation k and year t, w_{ikt}, is given by

$$w_{ikt} = \pi_{kt} + \alpha_{ik} + \beta_k \boldsymbol{s}_i + g_k(x_{ikt} - x^*) + \varepsilon_{ikt} \,, \tag{1}$$

where π_{kt} is a potentially time-varying occupation-specific wage premium; α_{ik} is an unobserved worker-occupation fixed effect; $\mathbf{s}_i$ is a vector of observable

[3] Using unusually rich data, Böhm et al. (forthcoming) are able to estimate across-occupation experience profiles, that is, the extent to which a year of work experience in one occupation increases the worker's productivity in this and all other occupations. In contrast, we estimate how wage growth in each occupation and at each point in time varies with overall potential labor market experience (given the limited length of our panel, we cannot construct workers' occupational histories). Therefore, our estimates have a different structural interpretation from those in Böhm et al. (forthcoming).

skills with its associated occupation-specific returns β_k; x_{ikt} is the worker's experience in the occupation measured in years and centered around x^*, to be discussed below; g_k is an occupation-specific experience profile; and ε_{ikt} is an i.i.d. shock. Our main goal is to estimate π_{kt} for each occupation, or at least its *change relative to a reference occupation*.

For the moment, let us assume that workers choose the occupation in which they earn the highest wage in each period, abstracting from dynamic considerations. Furthermore, let us assume for now that the shock ε_{ikt} is realized *after* workers have made their choice. These assumptions are the same as in Cortes (2016). This leaves us with two potential threats to identification: Selection on unobserved time-invariant characteristics, and occupation-specific experience profiles. We address these in turn.

2.1.1 Selection on Time-Invariant Characteristics

Consider the first difference of equation (1),

$$\Delta w_{ik} = \Delta\pi_k + g_k(x_{ikt} - x^*) - g_k(x_{ikt} - 1 - x^*) + \Delta\varepsilon_{ik}, \qquad (2)$$

where Δ is the first difference operator, so that $\Delta X = X_t - X_{t-1}$. If we estimate equation (2) using the sample of occupation stayers, we can be sure that selection on time-invariant skills α_{ik} and $\mathbf{s}_i$ is accounted for, since these terms are differenced out. An alternative method accomplishing this is of course to estimate equation (1) in levels and to include worker-by-occupation fixed effects, as in Cortes (2016). We prefer the first difference specification for two reasons. First, it allows us to run separate regressions for each year, and thus work with datasets of manageable size. Second, our data on wages and occupations come from repeated cross-sectional samples, so that it is difficult to construct long panels of individual workers, and to accurately capture longer occupational spells (see Section 3).

2.1.2 Occupation-Specific Experience Profiles

For concreteness, we approximate the profile by a polynomial of order M, $g_k(x) = \sum_{m=1}^{M} \gamma_{km}(x - x^*)^m$. Under this assumption, the component of wage growth due to experience – among occupation stayers – now becomes

$$g_k(x_{ikt} - x^*) - g_k(x_{ikt} - 1 - x^*) = \gamma_{k1} + \sum_{m=2}^{M} \gamma_{km}\{(x_{ikt} - x^*)^m - (x_{ikt} - 1 - x^*)^m\}.$$

The wage growth equation to be estimated is thus

$$\Delta w_{ik} = \Delta\pi_k + \gamma_{k1} + \sum_{m=2}^{M} \gamma_{km}\{(x_{ikt} - x^*)^m - (x_{ikt} - 1 - x^*)^m\} + \Delta\varepsilon_{ik}. \qquad (3)$$

Estimation of equation (3) for a given occupation yields a constant term $\theta_{kt} = \Delta\pi_k + \gamma_{k1}$. Thus, the challenge is to separate out changes in premia from the constant term of the experience profile. Note that γ_{k1} is the effect of additional experience at the point $x_{it} = x^*$. Human capital theory (Ben-Porath, 1967; Heckman et al., 1998) suggests that there comes a point in a worker's life cycle when human capital accumulation stops, or even reverses due to depreciation – a so-called flat spot where the marginal effect of experience on wages is zero. Thus, if x^* is set to be at the flat spot, then $\gamma_{k1} = 0$, solving the identification problem as we now have $\theta_{kt} = \Delta\pi_k$.[4]

We illustrate this strategy using a concrete example: The wage growth of *physical and engineering science technicians* from 2005–06. Figure 1 plots changes in log wages, together with the fitted polynomial, against potential experience re-centered around different values – the assumed locations of the flat spot. The fitted polynomial comes from estimating equation (3) choosing $m = 4$. Grey dashed lines mark the constant term estimated by the regressions, equal to premium growth under the assumption $\gamma_{k1} = 0$. The data reveal a strong downward trend in wage growth, consistent with faster skill accumulation among inexperienced workers, as well as a flattening of this relationship at higher levels of potential experience. The top-left panel does not re-center the data, thus yielding a large estimated premium growth of around 8 percent. But an assumption of zero skill accumulation for labor market entrants is of course highly implausible. Assuming flat spots at higher values such as 25, 30, or 35 all yield estimated premium growth around 2 percent, as shown in the remaining panels.

Figure 1 illustrates that choosing the flat spot means picking a point on the fitted first-differenced experience profile and attributing all wage growth at that point to growth in the premium.[5] Relying on a parametric prediction for the profile yields greater statistical power compared to simply using average wage growth at the flat spot.

[4] Our approach is related to Fosse and Winship (2019), who address the identification problem arising in the presence of age, cohort, and time effects. They highlight that it is only linear effects that are unidentified, and explain how one can bound these. However, a single restriction is often sufficient for point identification, as is the case in our context.

[5] To be precise, the flat spot assumption says that $g'_k(x^*) = 0$. In the polynomial case, $g'_k(x^*) = \sum_{m=1}^{M} \gamma_{km} m(x - x^*)^{m-1} = \gamma_{k1}$. Here, the flat spot assumption that $\gamma_{k1} = 0$ does not imply that that $\Delta w_{ik} \big|_{x=x^*} = \Delta\pi_k$ exactly, which requires $\sum_{m=2}^{M} \gamma_{km}\{-(-1)^m\} = 0$. However, in practice these equations will hold approximately, as is the case in Figure 1.

Figure 1: *Illustration of flat spot identification.*

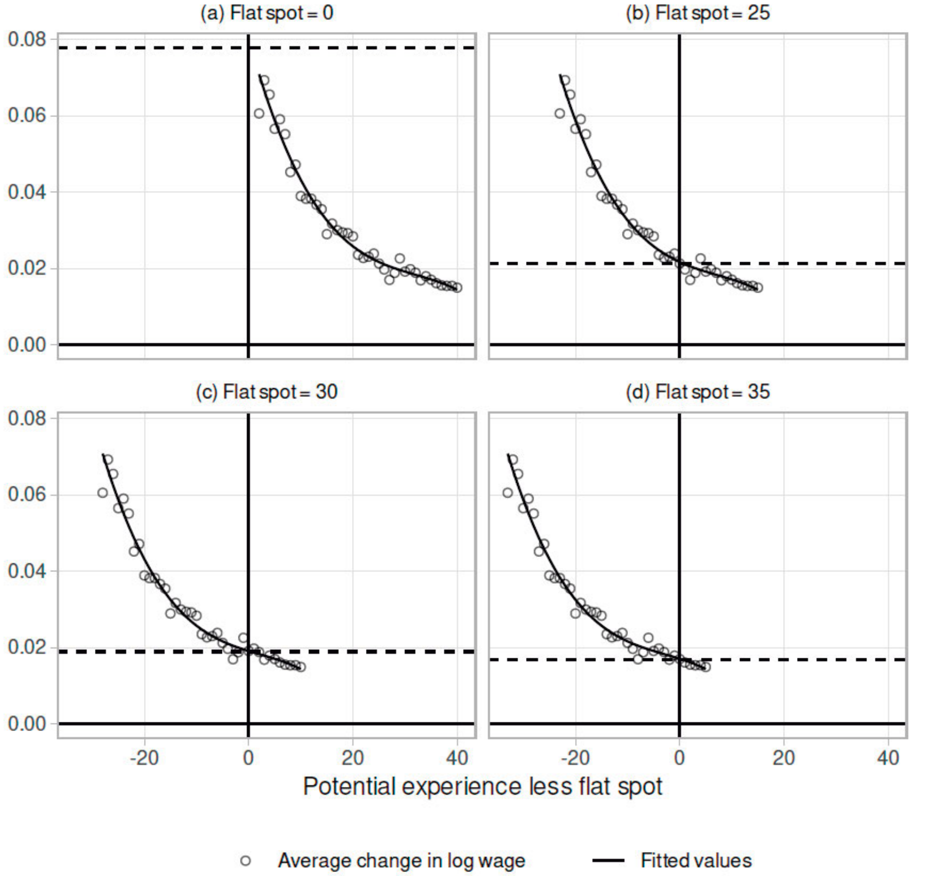

Notes: Grey dashed lines mark the constant term from estimating the experience profiles, equal to wage premium growth under the respective flat spot assumptions. The data include all individuals who worked as physical and engineering science technicians in 2005 and 2006. See Section 3 for further details on sample selection.

Figure 1 also raises the question whether the flat spot can be determined in a data-driven way. In general, the answer is no. Consider three hypothetical experience-wage profiles plotted in the top row of Figure 2. As we do not observe workers' time-invariant occupation-specific skills, we cannot estimate the profiles in levels. We thus first-difference the profiles, shown in the bottom row. The challenge remains to separate premium growth from skill accumulation. Consider first column (a). The differenced profile reproduces the nearly flat region of the original in-levels profile. While it may not be easy to determine the *exact* location of the flat spot, this would also not matter greatly for the estimated premium growth. However, recall that the econometrician cannot see the top row. As column (b) shows, a flat region in first differences can also result from a locally log-linear profile in levels. In this case, the true flat spot at 34 cannot be detected based on first differences. Finally, consider column (c), which shows a profile of roughly constant curvature and hence no flat region. Again, it is not obvious how to choose the flat spot based on the first-differenced profile.

Given that the flat spot cannot be identified without further assumptions, our approach is to set it at 30 for all occupations, while also reporting results for alternative values. In a further robustness check, we estimate flat spots under the additional assumption that the true profiles are strictly concave except for possible flat sections (that is, linear segments with non-zero slope, as in the middle column of Figure 2, are prohibited). In this case, the second derivative of the profile will be maximized (closest to zero) at the flat spot, so that any statistic of interest should change by the least amount – in absolute value – at the true flat spot. See Appendix A for further details.

A key advantage of our method is that it allows us to jointly estimate experience profiles and premium growth. Moreover, as we estimate separate models for each year, we essentially estimate time-varying experience profiles.[6] A third advantage is that we retain greater statistical power than existing approaches in the literature which implement the flat spot idea using only data on workers near the flat spot (Bowlus and Robinson, 2012; Cavaglia and Etheridge, 2020).[7]

We note that the interpretation of our estimated profiles is affected by the measurement of occupation-specific experience. With a panel that is relatively short (20 years) relative to the typical length of working lives, it is not possible to construct complete occupational histories for each worker.[8] In our baseline specification we therefore use potential overall labor market experience, based on age and years of schooling. Given this, the occupation-specificity of the γ_{km}'s means that experience is differently valued across occupations, but it does not matter in which occupation this experience was gained. Alternatively, one can simply interpret the estimated profiles as describing the wage growth in a given occupation and year as a function of potential overall labor market experience. This function will depend not only on deep structural parameters, but also on the characteristics – such as occupational histories – of the workers staying in that occupation in that year (and the year before).

[6] Strictly speaking, the experience profile in *levels* must be constant across the two adjacent years. This would not matter if we had specified the profile in changes in the first place. However, starting with a levels specification is arguably more natural given the Roy framework.

[7] It is also possible to implement flat spot identification via an iterative procedure (Lagakos et al., 2017).

[8] Another challenge is that private sector workers in Sweden are sampled, as we discuss in the data section.

Figure 2: *Simulated wage-experience profiles*

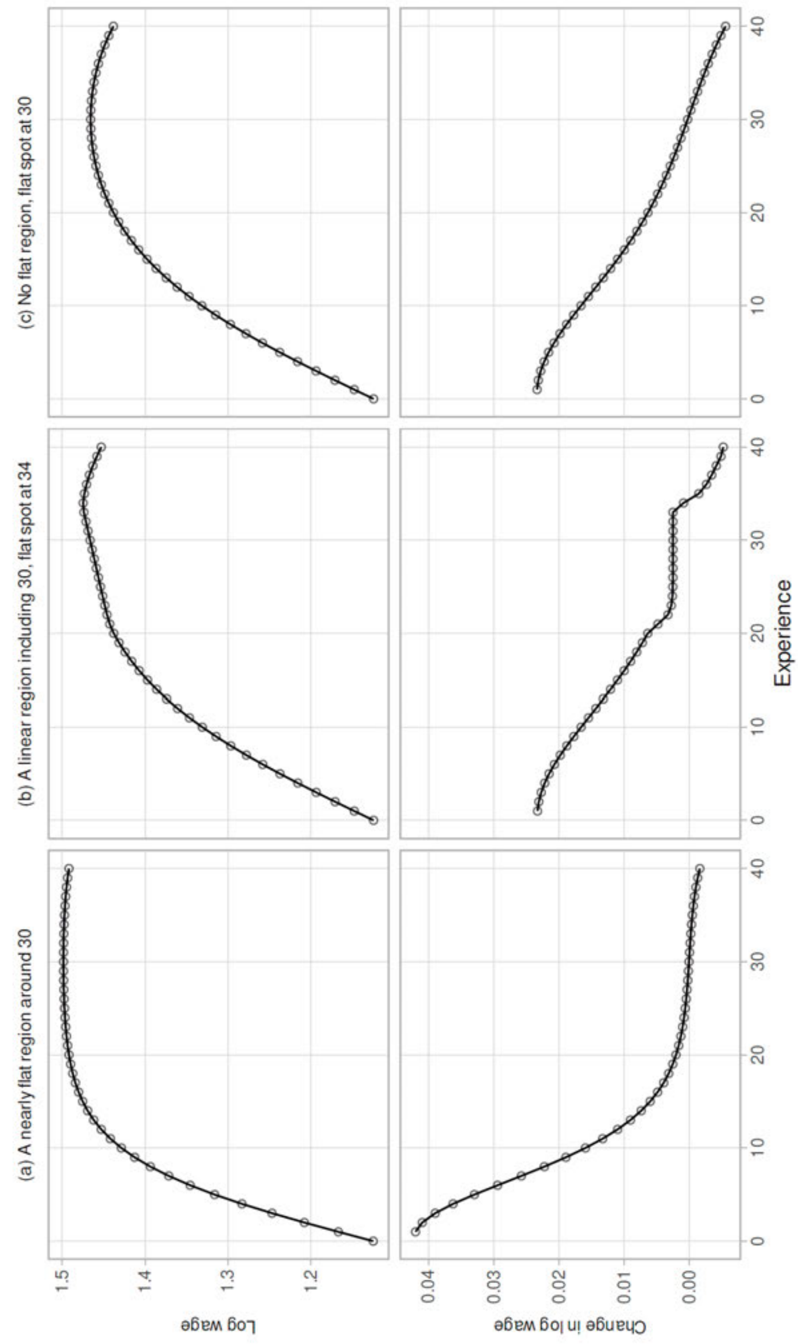

Notes: The top row shows various simulated experience profiles. The bottom row plots first differences of the profiles above, assuming wage premium growth of two percent.

2.1.3 Time-Varying Skill Returns

A key finding in recent research on inequality is that wage returns to various skills have evolved differently over time, with occupations seemingly playing an important role. While this is interesting in its own right, here we are mainly concerned with the impact of such changes on our ability to estimate changes in occupational wage premia. Suppose, then, that returns to portable skills vary over time,

$$w_{ikt} = \pi_{kt} + \alpha_{ik} + \beta_{kt}\boldsymbol{s}_i + g_k(x_{ikt} - x^*) + \varepsilon_{ikt},$$

so that wage growth now becomes

$$\Delta w_{ik} = \Delta\pi_k + (\Delta\beta_{kt})\boldsymbol{s}_i + \sum_{m=2}^{M} \gamma_{km}\{(x_{ikt} - x^*)^m - (x_{ikt} - 1 - x^*)^m\} + \Delta\varepsilon_{ik}. \quad (4)$$

For selected cohorts of Swedish men we actually have at our disposal the skill measures for which changing wage returns have been documented. We can thus assess whether our baseline estimates of $\Delta\pi_k$ are robust to controlling for these measures, by estimating equation (4) where the vector $\boldsymbol{s}_i$ contains cognitive and psycho-social skills, as described further in the data section.

2.1.4 Selection on Idiosyncratic Shocks

Let us now allow for selection on the idiosyncratic shock ε_{ikt}, as well. The constant term from estimating equation (3), imposing the flat spot assumption $\gamma_{k1} = 0$, now becomes $\theta_{kt} = \Delta\pi_k + E\left[\Delta\varepsilon_{ik} \mid k_{it} = k_{i,t-1} = k\right]$. The second term no longer equals zero, due to selection. Other things equal, occupations experiencing relatively fast premium growth will retain more workers with a bad realization of the shock, while occupations in which premia decline only retain those workers with very good realizations. Therefore, selection on idiosyncratic shocks biases downward the between-occupation variance in premium growth. This bias is more severe the larger is the variance of ε_{ikt}. A method to correct for this bias, developed by Böhm et al. (forthcoming), is to include occupation switchers in a regression of wage growth on workers' average choices. We implement this method as a robustness check.

2.1.5 Remaining Issues

There are a number of issues which are beyond the scope of this paper. These include forward-looking occupational choice, amenities, search frictions, and long-term wage contracts. We believe that addressing any one of these requires estimation of a fully specified structural model (for recent examples, see Roys and Taber, 2019; Traiberman, 2019).

2.2 Occupational Drivers of Changes in Wage Inequality

A key objective of this paper is to assess the importance of occupations for changes in wage inequality. Therefore, we need to formally characterize how changes in inequality relate to occupation-level changes such as differential premium growth and worker re-allocation. We closely follow Böhm et al. (forthcoming).

First, by the Law of Total Variance, $\mathrm{Var}(w_{it}) = \mathrm{E}[\mathrm{Var}(w_{it} \mid k)] + \mathrm{Var}(\mathrm{E}[(w_{it} \mid k)])$. That is, overall wage inequality can be decomposed into a within-occupation and a between-occupation component. Without specifying the distribution of skills, it is difficult to say much about how changes in premia affect the within component, so we focus on the between component.

To ease notation, let us from now on write $w_{kt} \equiv \mathrm{E}[w_{it} \mid k]$ and $\Delta w_k \equiv \Delta\mathrm{E}[\omega_i \mid k]$, and similarly for other variables. The difference operator $\Delta X \equiv X_1 - X_0$ denotes changes between two points in time 0 and 1, not necessarily adjacent years.

Note, to integrate out the conditioning variable – occupational choice – we must specify a distribution of occupational employment. When decomposing the variance at a given point in time, the obvious choice is to use the distribution at that point. But when considering changes over time, we need to be explicit about the distribution. We use subscripts to do so.

The change in between-occupation wage inequality can be written as

$$\mathrm{Var}_1(w_{k1}) - \mathrm{Var}_0(w_{k0}) = \underbrace{\mathrm{Var}_0(w_{k1}) - \mathrm{Var}_0(w_{k0})}_{\text{change at initial employment}} + \underbrace{\mathrm{Var}_1(w_{k1}) - \mathrm{Var}_0(w_{k1})}_{\text{re-allocation}}. \tag{5}$$

Define $y_{kt} \equiv w_{kt} - \pi_{kt}$, which captures workers' skills in the broadest sense – all parts of log wages not determined by the occupation premium. The first component on the right-hand side of equation (5) can be broken down as

$$\begin{aligned}\mathrm{Var}_0(w_{k1}) - \mathrm{Var}_0(w_{k0}) &= \mathrm{Var}_0(\Delta w_k) + 2\mathrm{Cov}_0(w_{k0}, \Delta w_k)\\ &= \mathrm{Var}_0(\Delta\pi_k) + \mathrm{Var}_0(\Delta y_k) + 2\mathrm{Cov}_0(\Delta\pi_k, \Delta y_k)\\ &\quad +2\mathrm{Cov}_0(w_{k0}, \Delta\pi_k) + 2\mathrm{Cov}_0(w_{k0}, \Delta y_k).\end{aligned} \tag{6}$$

From equation (6), we see how differential changes in premia may affect changes in wage inequality, and at the same time, how their effects may be offset by opposing forces. In particular, all the components of the decomposition involving changes in average skills Δy_k, as well as the re-allocation term from equation (5), can be seen as potentially countervailing effects due to workers' re-sorting. In contrast, all terms only involving $\Delta\pi_k$ and initial mean wages w_{k0} can be interpreted as giving the counterfactual increase in between-occupation wage inequality in the absence of re-sorting. That is, with worker composition unchanged, we have

$$\mathrm{Var}_0(w_{k1}) - \mathrm{Var}_0(w_{k0}) = \mathrm{Var}_0(\Delta\pi_k) + 2\mathrm{Cov}_0(w_{k0}, \Delta\pi_k), \tag{7}$$

which is a key object of interest in our analysis. Equation (7) shows that, holding worker composition constant, changes in wage premia have a large effect on wage inequality if they are very dispersed, or if they are positively correlated with initial mean wages.

3. Data Description

3.1 Data Sources

We obtain demographic information (year of birth, sex, municipality of residence, education, immigration status) from Statistics Sweden's LISA database, covering the population of Swedish residents in 1985-2016. LISA also contains employment status in November each year, annual salary income, as well as industry and municipality of workplace.

Some information that is key for our purposes is absent from LISA. In particular, LISA does not contain weeks and hours worked, nor occupation. For this, we turn to a database called Swedish Wage Structure Statistics (henceforth WSS). WSS contains three-digit occupation codes according to the *SSYK96* classification for 1996–2013, and according to the *SSYK2012* classification for 2014-2016.[1] The two classifications cannot be mapped unambiguously, and breaks in employment trends are apparent even at higher levels of aggregation. We therefore end our main analysis in 2013.

WSS also contains contractual monthly wage rates. This in combination with annual salary income allows us to determine annual labor supply. Most importantly, these contractual wage rates are the main outcome of interest for our analysis, since we are interested in the price of labor.

A drawback of WSS is that outside the public sector, only a sample of workers is available. Sampling is stratified by firms, with large firms being more likely to be drawn. This does not pose any problems for cross-sectional analysis – sampling weights are provided – but makes it more difficult to analyze dynamic phenomena such as occupational mobility. We discuss this issue further in the next sub-section.

For some of our analysis, we use test scores collected during military enlistment in the last decades of the 20th century, after which conscription was gradually phased out. Among the 1952-1981 birth cohorts, more than 90 percent of Swedish-born males are covered by these data. We use a combined measure of cognitive skills based on four different standardized tests of inductive, verbal and spatial skills, and technical comprehension, and a measure of psycho-social skills (sometimes called "non-cognitive skills") based on a half-

[1] SSYK stands for *Standard för Svensk Yrkesklassificering*, literally "Standard for Swedish occupation classification", a version of the International Standard Classification of Occupations (ISCO).

hour, semi-structured interview with a certified psychologist.[2] We standardize the two measures within each draft cohort to have a mean of zero and a standard deviation of one. To ensure comparability, we estimate our main specification also for the sub-sample of male cohorts for which enlistment data are available.

3.2 Sample Selection and Construction of Variables

Our population of interest includes all Swedish employees aged 18–64 during the years 1996–2013 (sometimes extended to 2016). Employees are individuals who are employed in November and whose annual labor earnings are no less than three times the 10th-percentile monthly wage. We calculate individual wage growth for all adjacent years, dropping anyone with wage growth below the first or above the 99th percentile for each pair of years.[3]

We calculate potential labor market experience as years elapsed since year of graduation, based on highest education attained and a school starting age of six. To reduce noise, we drop observations with potential experience below two and greater than 40 years. Due to the limited length of the panel as well as due to sampling, we are unable to construct actual occupation-specific experience.

We use sampling weights to adjust for stratification. The raw weights supplied in WSS feature some extremely large values, and this may introduce noise, especially when multiplying the weights for a first-difference analysis using a two-year panel. Whenever we work with individual, two-year panel data, we therefore trim the weights following the procedure of Potter (1990).[4] However, when computing aggregate moments, we use the original weights.

For our baseline analysis we use the 3-digit-level *SSYK96* occupational classification, which includes 101 occupations. However, we sometimes use a coarser classification for descriptive and other purposes.

[2] The intent of the interview was to evaluate the psychological fitness for coping with military service. See Lindqvist and Vestman (2011) and Fredriksson et al. (2018) for more details on these data.

[3] Extreme values of wage growth – five or more standard deviations away from the mean – may occur because individuals enter into and exit from executive positions (Skans et al., 2009). We drop extreme values as these can have a large impact on the results.

[4] The procedure is as follows. We first fit a Beta(α,β) distribution to the weights. Second, weights whose estimated cumulative probability is above 99 percent are trimmed to the estimated 99th percentile. Third, weights are re-scaled such that their sum is unchanged. This procedure is repeated ten times.

Figure 3: *Growth in wages, premia, and skills against employment growth, 1996–2013*

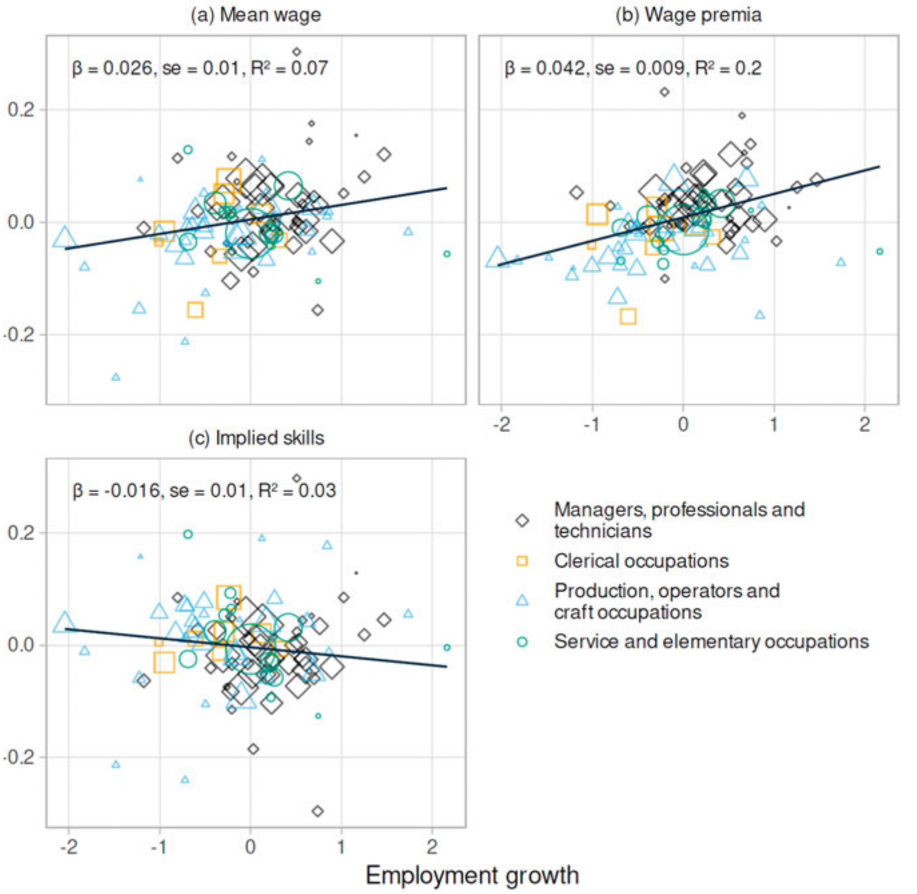

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against the change in log employment. Wage premia are estimated according to our baseline specification equation (8). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

4. Results

4.1 Raw Wages, Wage Premia, and Employment

To set the stage, we document the relationship between growth in averages wages and growth in employment as well as initial wages, across occupations for the period 1996–2013. Panel (a) of Figure 3 plots the long difference in log wages against the long difference in the log of employment, with each marker representing one occupation. First, by moving along the horizontal axis, we see much variation in employment growth. Production, operators, and craft occupations tend to see low (often negative) employment growth, while on average, employment growth appears highest among managers, professionals, and technicians. Clerical and services occupations fall somewhere in between. However, there is much variation even within these broad categories. Turning to wage growth, there is a positive but rather weak relationship with employment growth. Panel (a) of Figure 4 reveals an even weaker relationship of average wage growth with initial (1996) average wages.

However, as discussed in Section 2, average wage growth captures both changes in occupational wage premia and changes in worker composition and hence average skills. In order to isolate changes in wage premia, we operationalize equation (2) by estimating separate regressions of year-on-year changes in individual log wages on occupation fixed effects and a polynomial in potential experience:

$$\Delta w_{it} = \varphi_{kt} + \sum_{m=2}^{4} \gamma_{km} \{\tilde{x}_{it}^{m} - (\tilde{x}_{it} - 1)^{m}\} + u_{it} \tag{8}$$

where φ_{kt} are occupation-specific fixed effects; $\tilde{x}_{it} \equiv x_{it} - x^{*}$ is potential experience re-centered around the assumed flat spot in the experience profile; and γ_{km} are polynomial coefficients allowed to vary by occupation. In our main specification, we use a fourth-order polynomial and re-center potential experience at 30 years. We report robustness checks with respect to these choices below. We estimate separate regressions for each pair of adjacent years in our sample. In order to control for changes in worker composition, we use only individuals who remained in the same occupation across both years, $k_{it} = k_{i,t-1}$.

Under the assumption that there is no selection on idiosyncratic shocks and that $\gamma_{m1} = 0$ (the flat spot assumption), the fixed effects φ_{kt} consistently estimate premium growth $\Delta_{t-1}^{t}\pi_k$ for an adjacent pair of years. We estimate premium growth over the full period by simply accumulating the estimated year-on-year changes, $\widehat{\Delta_{1997}^{2013}}\pi_k = \sum_{t=1997}^{2013} \hat{\varphi}_{kt}$.

Our premium growth estimates are plotted against employment growth in Panel (b) of Figure 3. The relation between premium growth and employment growth is stronger than that of mean wage growth – the slope is steeper, and R^2 almost triples. This pattern implies that while demand factors were pushing up wage premia during this period, changes in the skill composition of workers acted as a counteracting force, resulting in the tempered trend we see in average wage growth. This is consistent with a situation where growing labor demand in certain occupations attracts new workers with lower productivity than the incumbents – and conversely, occupations with falling labor demand might let their lower-productivity workers go first. The implied change in skill composition can be backed out from our estimates by simply subtracting the estimated changes in premia from the observed changes in average wages. This is shown in Panel (c) of Figure 3. As expected, faster growing occupations have seen falling implied skill levels in their workforce, although this relationship is not very strong.

Panel (b) of Figure 4 shows that premium growth is strongly positively associated with initial wages. Given equation (6), this suggests that premium growth would cause an increase in between-occupation wage inequality in the absence of compositional changes. However, panel (c) of Figure 4 already gives an idea of how strong these compositional changes might be – growth in average skills are strongly negatively related to initial wages. We explore these issues in detail in Section 4.2.

One way to assess the plausibility of the estimated growth in skills is to check its association with changes in years of schooling. Panel (c) Figure 5 shows that there is indeed a positive relationship, with a fairly high R^2 of 0.2. On the other hand, panel (b) of the same figure shows a negative association between premium growth and changes in years of schooling, consistent with lower educated workers moving into occupations experiencing positive demand changes.[5]

While the evidence presented so far suggests that the forces predicted by the Roy model are at work, it remains to assess their quantitative importance for the evolution of wage inequality in Sweden. We do so next.

[5] For completeness, Figure A2 displays the respective bi-variate associations of wage growth, premium growth, and implied skill growth, showing positive correlations between wage growth and premium growth, and wage growth and skill growth, and a negative correlation between premium growth and skill growth.

Figure 4: *Growth in wages, premia, and skills against initial wages, 1996–2013*

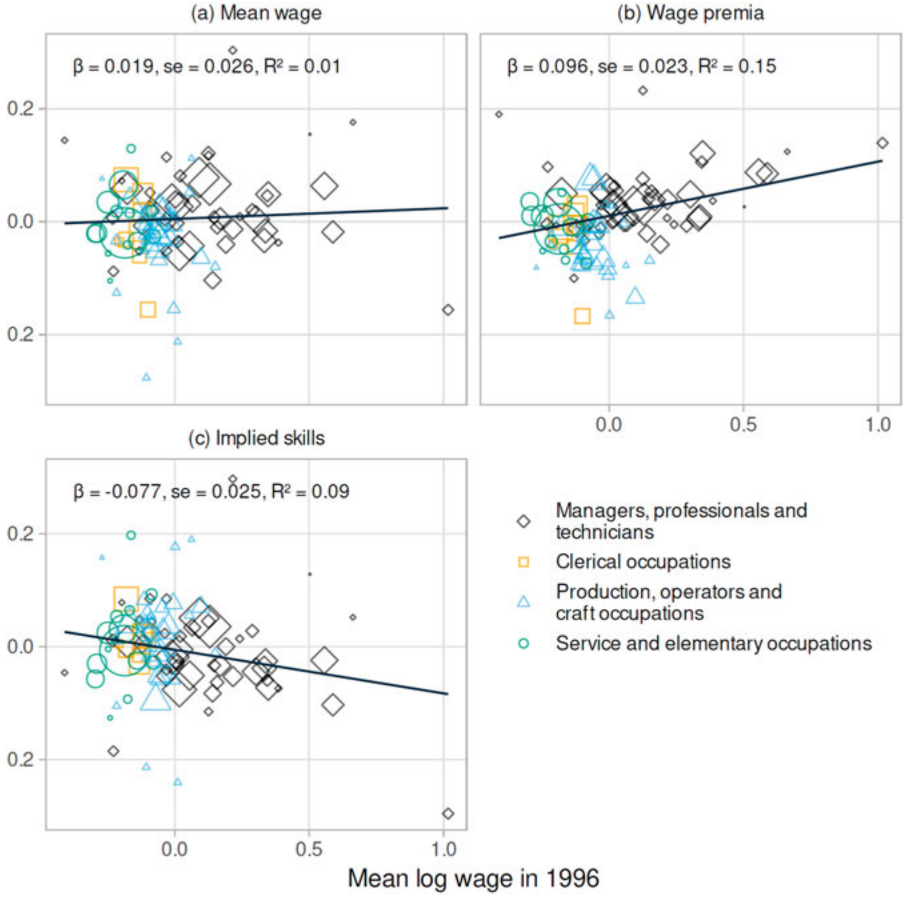

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against initial mean log wages. Wage premia are estimated according to our baseline specification equation (8). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

Figure 5: *Growth in wages, premia, and skills against growth in schooling, 1996–2013*

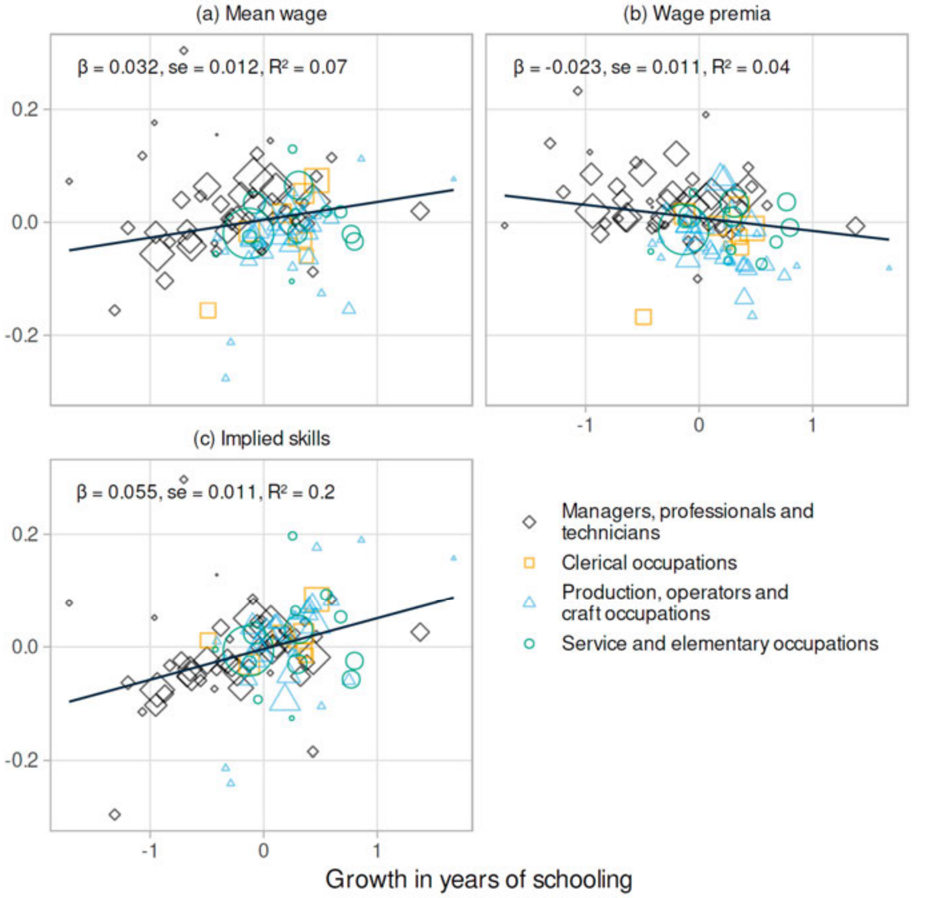

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against initial the growth in average years of schooling. Wage premia are estimated according to our baseline specification equation (8). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

4.2 Decomposing Changes in Between-Occupation Wage Inequality

To quantify the role of differential premium growth for changes in between-occupation inequality in Sweden, we use our estimates to calculate the counterfactual scenarios developed in Section 2.2. We first focus on the long difference 1996–2013 and then examine changes at annual frequency.

The first three lines in column (1) of Table 1 show the change in the observed variance of log wages, the change in between-occupation variance, as well as the change in betweenoccupation variance holding occupational employment fixed at 1996. The variance of log wages increased by 0.026 in 1996–2013, from 0.073 in 1996 (to avoid excessive decimal places, we multiply the variance and its components by 100 from here on). Although the wage distribution in Sweden is still highly compressed compared to other countries (Graetz, 2020), this increase is large in relative terms.

Table 1: *Decomposition of changes in between-occupation wage inequality*

	(1)	(2)	(3)	(4)
	Baseline	Common flat spot		Occ.-spec.
		25	35	flat spot
Total				
$\Delta\text{Var}(w_{ik})$	2.57			
Between				
$\Delta\text{Var}(w_k)$	1.31			
$\Delta\text{Var}_0(w_k)$	.39			
Components				
$\text{Var}_0(\Delta\pi_k) + 2\text{Cov}_0(w_{k0}, \Delta\pi_k)$	.94	2.03	.29	.66
$\text{Var}_0(\Delta\pi_k)$	.23	.43	.17	.27
$2\text{Cov}_0(w_{k0}, \Delta\pi_k)$	.71	1.59	.12	.4
$\text{Var}_0(\Delta y_k)$	.26	.37	.24	.25
$2\text{Cov}_0(w_{k0}, \Delta y_k)$	-.57	-1.45	.02	-.26
$2\text{Cov}_0(\Delta\pi_k, \Delta y_k)$	-.24	-.55	-.16	-.27

Notes: The table reports results from a decomposition of the change in the between-occupation variance in wages between 1996 and 2013 for different flat spot levels. See equation (6) for the formal statement of the decomposition. Column (1) uses a common flat spot for all occupations, at 30 years of potential experience, when estimating growth in wage premia. Columns (2)–(4) vary this common flat spot as indicated. Column (5) estimates a flat spot for each occupation using the procedure described in Appendix A. All figures have been multiplied by 100 for readability.

Between-occupation wage inequality accounts for just over half of the increase in overall variance. But this is allowing for the employment weights in the calculation of variance to change over time. If employment shifts from middle- to both high- and low-paying occupations, we should expect between-occupation inequality to increase even if wage premia do not grow differentially. The phenomenon of job polarization has been extensively documented in the literature (Goos et al., 2014; Adermon and Gustavsson, 2015), and Figure A1 confirms that it is present also in our sample period.

Our main interest, however, is in occupation-level drivers of wage inequality that are due to differential changes in compensation for a fixed set of workers. The third row in Table 1 shows that holding employment fixed at 1996 levels, the contribution of between-occupation variance shrinks by more than two thirds. But, as discussed above, changes in observed wages at the occupation level may mask changes in composition. To assess the role of differential growth in occupational wage premia, we perform the decomposition given by equation (6).

Column (1) of Table 1 presents our baseline results, with the flat spot set at 30 for all occupations. Holding worker composition constant, the increase in between-occupation variance would have been 0.94 based on our decomposition. This is more than twice the increase in the between-occupation variance of raw wages (at constant employment), and almost 40 percent of the increase in the overall variance of log wages. Most of this effect is due to a positive covariance between initial wages and premium growth, while the variance in premium growth plays a relatively minor role. The last two rows in column (1) of Table 1 show the attenuating forces: Changes in worker skills are negatively correlated with both initial wages and growth in wage premia.

Figure 6 shows the evolution of the variance components year-on-year. Interestingly, during the period 1996–2001, which saw the fastest growth in wage inequality, the attenuating forces of a changing skill composition are absent, and the no-sorting counterfactual closely tracks between-occupation inequality in raw wages (at constant employment). The attenuating forces emerge only after 2001.[6]

4.3 Robustness Checks

We conduct a number of robustness checks for the results that depend on the estimation of wage premium growth. First, we vary the location of the flat spot. As expected given the shape of wage-experience profiles and the above discussion of Figure 1, the decomposition results are sensitive to the choice of flat spot, as seen in columns (3) and (4) of Table 1. The sensitivity varies by component: The variance of premium growth appears more stable than the covariance of premium growth and initial wages.[7]

[6] Column (1) in panel B of Table A1 displays the decomposition results for the sub-period 2001–2013. Figures A3 and A4 display the relationships between growth in wages, premia, and implied skills on the one hand, and employment growth and initial wages on the other, for 2001–2013. While overall inequality changed little during this time, the pattern of premium growth and compositional changes is qualitatively very similar to that for the whole sample period.

[7] As premium growth and skill growth are strongly negatively correlated, this difference in sensitivity is mirrored by the other components.

Figure 6: *Decomposition of changes in between-occupation inequality 1996-2013*

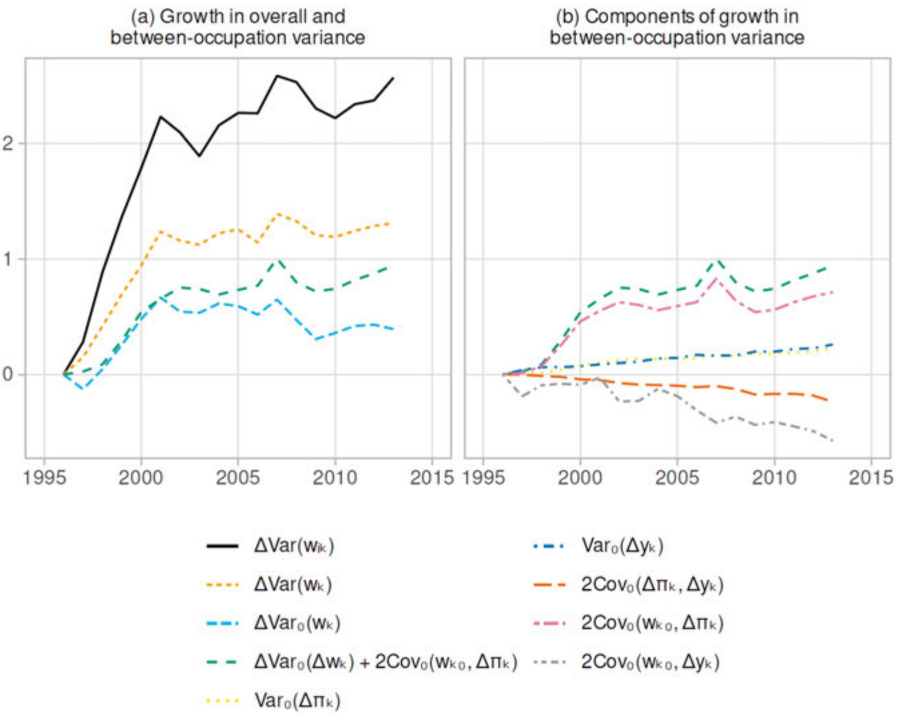

Notes: The figure plots the results from the decomposition given by equation (6) for every year pair $\{1996, t\}\ \forall\ t \in \{1996, \ldots, 2013\}$.

However, when we attempt to determine occupation-specific flat spots in a data-driven way – based on the assumption of strictly concave wage-experience profiles (except for possible flat regions) as discussed in Appendix A – we obtain results quite similar to our baseline specification (column 5). Note also that setting the flat spot at zero, which would be implied if we simply added higher-order terms of potential experience without re-centering them, yields clearly unreasonable results (column 2 of Table A2).

Table A2 displays the decomposition components for a set of further robustness checks. These include changing the order of the polynomial in potential experience; adjusting for endogenous mobility using the method of Böhm et al. (forthcoming); allowing for differential growth in wage premia at the level of regions and industries; pooling the data to estimate time-invariant experience profiles; restricting the data to men with non-missing enlistment scores; controlling for time-varying returns to cognitive and non-cognitive skills within this restricted sample; and dividing the data by gender. The results are robust in the sense that the no-sorting counterfactual in the majority of cases is of similar or even larger magnitude compared to the baseline.[8]

Finally, we probe the robustness of the associations of premium growth and implied skills with employment growth, initial wages, and years of schooling. The results are shown in Figures A5 and A6, and once again are largely similar to the baseline specification.

Figure 7: *Estimated occupational experience profiles for selected occupations and years*

Notes: The figure plots the estimated experience profiles from equation (8) for the indicated occupations and years.

4.4 Changes in Occupational Experience Profiles

A key advantage of our empirical approach is that we are able to estimate occupational experience profiles that vary over time. We estimate profiles for 101 occupations and each pair of years in 1996–2013. Due to space constraints, we only show estimated profiles for the largest (in terms of average

[8] Note that using a polynomial of order one or forcing the experience profiles to be constant over time are more restrictive and thus inferior to our baseline specification.

employment in 1996–2013) 3-digit occupation in each of nine main categories, for the years 1997, 2002, 2008, and 2013.

The estimated profiles are shown in Figure 7. There are several noteworthy findings. First, in all occupations wage growth is fastest for inexperienced workers, but this pattern is much more pronounced in some occupations (finance & sales professionals, building frame workers) than in others (personal care workers). Second, while in some occupations the profiles are stable (building frame workers), in others they show large changes over time (computing professionals). Third, profiles are steepest in the late 1990s in several cases, but this not a universal pattern.

To further investigate changes over time, we plot the median as well as quartiles of two measures capturing the steepness of the profiles, namely, the value of the profile at ten years of potential experience as well as the maximum value (both in levels). Panel (a) of Figure 8 reveals that, by both measures, profiles were indeed somewhat steeper in the late 1990s. But even more striking is that the steepness of the profiles was much more dispersed in that period.

Finally, we explore if there is a systematic relationship between dispersion in wage-experience profiles and dispersion in wage premium growth. Panel (b) of Figure 8 plots the variances of the two steepness measures along with the variance of premium growth against time. It appears that years with higher dispersion in profiles also tend to see higher dispersion in premium growth.

Figure 8: *Wage-experience profiles over time*

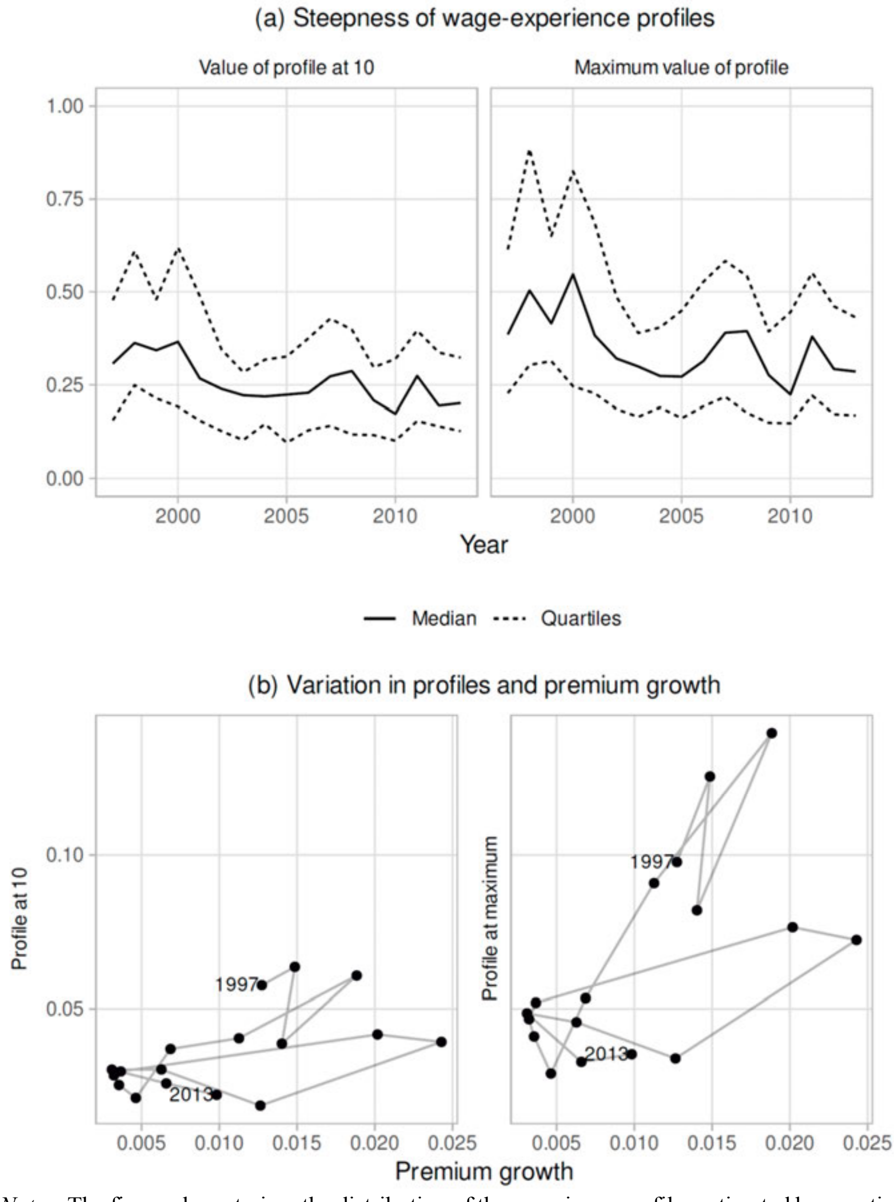

Notes: The figure characterizes the distribution of the experience profiles estimated by equation (8) over time (panel (a)) and shows how variance in selected characteristics of experience profiles is related to variance in premium growth (panel (b)).

5. Conclusion

We contribute to the literature on shifts in the wage structure by jointly estimating growth in occupational wage premia and occupation-specific life cycle wage profiles. We document substantial changes in occupations' relative premia in Sweden in recent decades, which are masked in the raw wage data due to worker sorting. There is a positive association between premium growth and employment growth, suggesting that workers have been responsive to changes in occupational demand. The relative premia changes are estimated to have substantially contributed to the increase in overall wage inequality. We also document large heterogeneity in life-cycle profiles across occupations, as well as substantial shifts of the profiles over time. Allowing for occupation-level changes in returns to cognitive and psycho-social skills has little effect on the results.

Our results suggest that although the overall wage structure in Sweden is highly compressed, forces related to technological change do influence the wage structure and drive workers' occupational choices. An open question is why the increase in Swedish wage inequality was concentrated in the late 1990s. This could be due to a temporary rise in the flexibility of collective bargaining, or it may reflect uneven technological change, for instance a transitional period of technology adoption (Beaudry et al., 2016).

The method we propose to estimate changes in occupational wage premia may fruitfully be applied to other settings, especially those in which experience profiles appear to change over time, and in cases where only short (two-year) panels of workers are available.

References

Acemoglu, D. and D. Autor (2011): “Skills, Tasks and Technologies: Implications for Employment and Earnings,” in *Handbook of Labor Economics*, ed. by O. Ashenfelter and D. Card, Elsevier, vol. 4, chap. 12, 1043–1171.

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

Autor, D. H., L. F. Katz, and M. S. Kearney (2006): “The Polarization of the U.S. Labor Market,” *American Economic Review*, 96, 189–194.

Barany, Z. L. and C. Siegel (2018): “Job Polarization and Structural Change,” *American Economic Journal: Macroeconomics*, 10, 57–89.

Beaudry, P., D. A. Green, and B. M. Sand (2016): “The Great Reversal in the Demand for Skill and Cognitive Tasks,” *Journal of Labor Economics*, 34, S199–S247.

Ben-Porath, Y. (1967): “The Production of Human Capital and the Life Cycle of Earnings,” *Journal of Political Economy*, 75, 352–365.

Böhm, M. J., H.-M. von Gaudecker, and F. Schran (forthcoming): “Occupation Growth, Skill Prices, and Wage Inequality,” *Journal of Labor Economics*.

Bowlus, A. J. and C. Robinson (2012): “Human Capital Prices, Productivity, and Growth,” *American Economic Review*, 102, 3483–3515.

Böhm, M. J. (2020): “The Price of Polarization: Estimating Task Prices under Routine-Biased Technical Change,” *Quantitative Economics*, 11, 761–799.

Cavaglia, C. and B. Etheridge (2020): “Job Polarization and the Declining Quality of Knowledge Workers: Evidence from the UK and Germany,” *Labour Economics*, 66, 101884.

Cortes, G. M. (2016): “Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data,” *Journal of Labor Economics*, 34, 63–105.

Deming, D. J. (2017): “The Growing Importance of Social Skills in the Labor Market,” *The Quarterly Journal of Economics*, 132, 1593–1640.

Deming, D. J. (2021): “The Growing Importance of Decision-Making on the Job,” Working Paper 28733, National Bureau of Economic Research.

Deming, D. J. and K. Noray (2020): “Earnings Dynamics, Changing Job Skills, and STEM Careers,” *The Quarterly Journal of Economics*, 135, 1965–2005.

Edin, P.-A., P. Fredriksson, M. Nybom, and B. Öckert (2022): “The Rising Return to Noncognitive Skill,” *American Economic Journal: Applied Economics*, 14, 78–100.

Fosse, E. and C. Winship (2019): “Bounding Analyses of Age-Period-Cohort Effects,” *Demography*, 56, 1975–2004.

Fredriksson, P., L. Hensvik, and O. N. Skans (2018): “Mismatch of Talent: Evidence on Match Quality, Entry Wages, and Job Mobility,” *American Economic Review*, 108, 3303–38.

Goos, M. and A. Manning (2007): “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain,” *The Review of Economics and Statistics*, 89, 118–133.

Goos, M., A. Manning, and A. Salomons (2014): “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 104, 2509–26.

Graetz, G. (2020): “Technological Change and the Swedish Labor Market,” Working Paper 2020:19, IFAU.

Heckman, J. J., L. Lochner, and C. Taber (1998): “Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents,” *Review of Economic Dynamics*, 1, 1–58.

Hoffmann, F., D. S. Lee, and T. Lemieux (2020): “Growing Income Inequality in the United States and Other Advanced Economies,” *Journal of Economic Perspectives*, 34, 52–78.
Lagakos, D., B. Moll, T. Porzio, N. Qian, and T. Schoellman (2017): “Life Cycle Wage Growth across Countries,” *Journal of Political Economy*, 126, 797–849.
Lindqvist, E. and R. Vestman (2011): “The Labor Market Returns to Cognitive and Noncognitive Ability: Evidence from the Swedish Enlistment,” *American Economic Journal: Applied Economics*, 3, 101–128.
Liu, D. C. and J. Nocedal (1989): “On the limited memory BFGS method for large scale optimization,” *Mathematical Programming*, 45, 503–528.
Potter, F. J. (1990): “A study of procedures to identify and trim extreme sampling weights,” in *Proceedings of the American Statistical Association, Section on Survey Research Methods*, American Statistical Association Washington, DC, vol. 225230.
Roy, A. D. (1951): “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers*, 3, 135–146.
Roys, N. A. and C. R. Taber (2019): “Skill Prices, Occupations, and Changes in the Wage Structure for Low Skilled Men,” Working Paper 26453, National Bureau of Economic Research.
Skans, O. N., P.-A. Edin, and B. Holmlund (2009): “Wage Dispersion Between and Within Plants: Sweden 1985-2000,” in *The Structure of Wages: An International Comparison*, National Bureau of Economic Research, Inc, NBER Chapters, 217–260.
Traiberman, S. (2019): “Occupations and Import Competition: Evidence from Denmark,” *American Economic Review*, 109, 4260–4301.
Wright, E. O. and R. E. Dwyer (2003): “The patterns of job expansions in the USA: a comparison of the 1960s and 1990s,” *Socio-Economic Review*, 1, 289–325.

Appendix A: Procedure for Estimating Occupation-Specific Flat Spots

Suppose that experience profiles are strictly concave except for possible flat regions. That is, linear segments with non-zero slope, as in the middle column of Figure 2, are prohibited. Formally, $g''(x) \leq 0, g''(x) = 0 \Rightarrow g'(x) = 0$. This implies that the second derivative of the profile will be maximized (closest to zero) at the flat spot, so that any statistic of interest should change by the least amount – in absolute value – at the true flat spot. We use this insight to pin down the flat spot in a data-driven way.

Recall from Section 2.2 that the change in between-occupation variance of log wages, at constant employment, can be decomposed as

$$\begin{aligned} \mathrm{Var}_0(w_{k1}) - \mathrm{Var}_0(w_{k0}) &= \mathrm{Var}_0(\Delta w_k) + 2\mathrm{Cov}_0(w_{k0}, \Delta w_k) \\ &= \mathrm{Var}_0(\Delta \pi_k) + \mathrm{Var}_0(\Delta y_k) + 2\mathrm{Cov}_0(\Delta \pi_k, \Delta y_k) \\ &\quad + 2\mathrm{Cov}_0(w_{k0}, \Delta \pi_k) + 2\mathrm{Cov}_0(w_{k0}, \Delta y_k). \end{aligned} \tag{A1}$$

Denote by μ the components of the decomposition,

$$\mu \in \mathcal{M} \equiv \{\mathrm{Var}_0(\Delta\pi_k), \mathrm{Var}_0(\Delta y_k), 2\mathrm{Cov}_0(w_{k0}, \Delta\pi_k), 2\mathrm{Cov}_0(\Delta\pi_k, \Delta y_k)\}.$$

Each of the elements of $\mathcal{M}$ depends on the change in the premia $\Delta\pi_k$, which in turn depend on the chosen flat spots. However, the sum of all components on the right-hand side of equation (A1) is constant, so we exclude $2\mathrm{Cov}_0(w_{k0}, \Delta y_k)$ from the set $\mathcal{M}$.

Let ϖ denote the vector of changes in premia, and let $\tilde{\mathbf{x}}$ denote the vector of candidate flat spots. Both vectors contain K elements, where K is the total number of occupations, indexed by k. We denote the above-mentioned functional dependence by $\mu \equiv \mu(\varpi(\tilde{\mathbf{x}}))$. Using the chain rule, we define the sensitivity of μ to changing the flat spot, in absolute terms, as

$$|d\mu(\varpi(\tilde{\mathbf{x}}))| \equiv \left| \sum_k \frac{\partial \mu}{\partial(\Delta\pi_k)} \times \sum_{k'} \frac{\partial(\Delta\pi_{k'})}{\partial \tilde{x}_{k'}} \times d\tilde{x}_{k'} \right|.$$

Under strictly concave experience profiles, we conjecture that $|d\mu(\varpi(\tilde{\mathbf{x}}))|$ attains its minimum at or near the vector of true flat spots $\mathbf{x}^*$, and similarly for the sum over $|d\mu(\varpi(\tilde{\mathbf{x}}))|$,

$$\mathbf{x}^* = \arg\min_{\tilde{\mathbf{x}}} \sum_{\mu \in \mathcal{M}} |d\mu(\varpi(\tilde{\mathbf{x}}))| \tag{A2}$$

We implement the optimization problem given by equation (A2) in practice by solving

$$\mathbf{x}^* = \arg\min_{\tilde{\mathbf{x}}} S \times \sum_{\widehat{\mu} \in \mathcal{M}} \left[(\widehat{\mu}(\tilde{\mathbf{x}} + \tau) - \widehat{\mu}(\tilde{\mathbf{x}}))^2 + (\widehat{\mu}(\tilde{\mathbf{x}} - \tau) - \widehat{\mu}(\tilde{\mathbf{x}}))^2 \right]$$

where $\hat{\mu}$ denote the estimated moments, τ is size-k vector with constant elements representing step size, and S is a scaling factor chosen for numerical stability. We set the elements of τ to equal 0.01 and S = 1e+7. We use the L-BFGS-B method (Liu and Nocedal, 1989) implemented by the optim package in R. We impose $\tilde{x}_k \in [25{,}40]\ \forall k$. As the procedure appears to be sensitive to initial values, we draw initial values at random from the continuous uniform distribution $U(26{,}39)$ for each $\tilde{x}_k$. This process is repeated 100 times. We then choose the $\mathbf{x}^*$ with the lowest associated loss.

Note that in principle, given strictly concave profiles one should be able to find the flat spots by minimizing the sensitivity of the $\Delta\pi_k$'s instead of a moment that is a function of them. However, approximating the experience profiles by a polynomial does not guarantee that the estimated profiles are actually strictly concave. Alternatively, one could impose a functional form on the profiles that does guarantee strict concavity. We attempted to do this, but the estimation turned out to be highly unstable.

Appendix B: Additional Figures and Tables

Figure A1: *Job polarization*

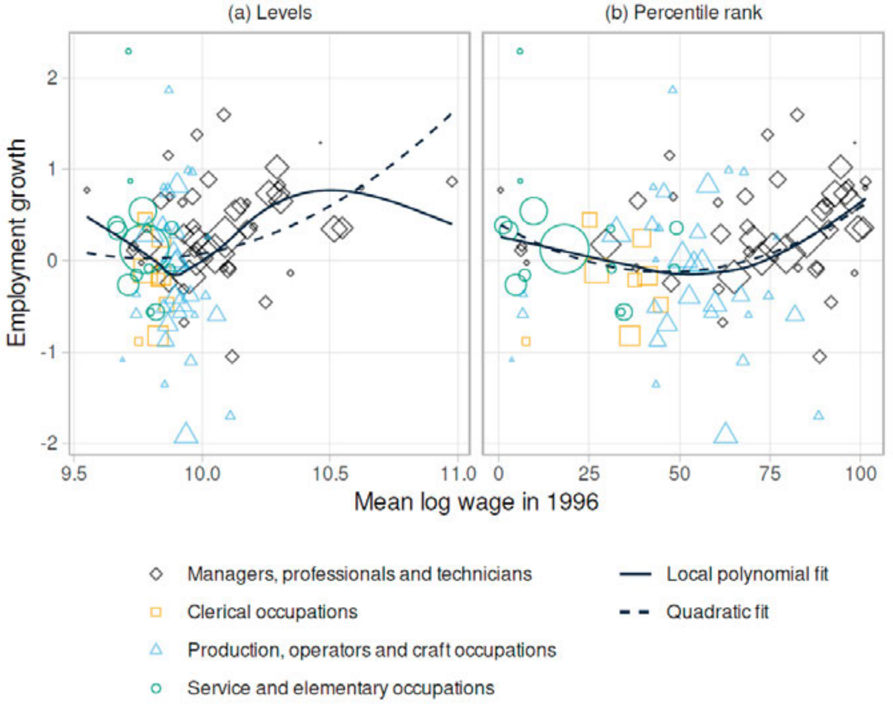

Notes: The figure plots the growth in log employment against mean log wages in 1996. In Panel (b), log wages have been percentile-ranked, weighted by initial employment. Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression lines are weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

Figure 2: *Relations between growth rates*

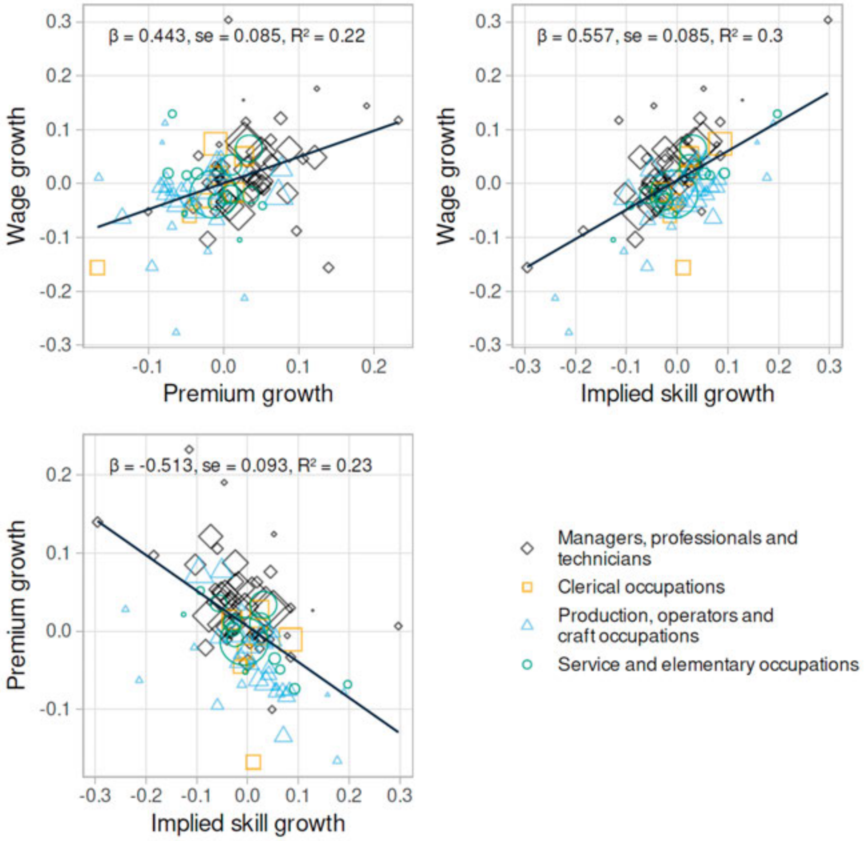

Notes: The figure plots the bivariate relationships between the growth in mean log wages, cumulative estimated wage premia, and the implied change in mean skills. Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression lines are weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

Figure A3: *Growth in wages, premia, and skills against employment growth, 2001–2013*

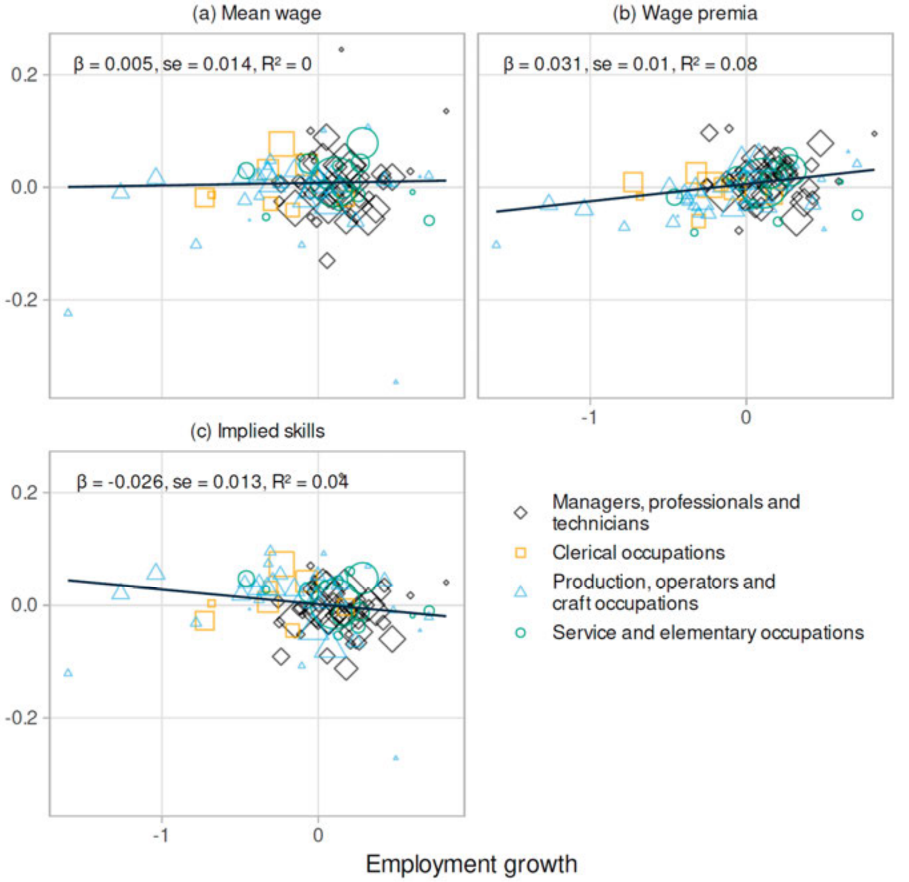

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against the change in log employment. Wage premia are estimated according to our baseline specification equation (8). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

Figure A4: *Growth in wages, premia, and skills against initial wages, 2001–2013*

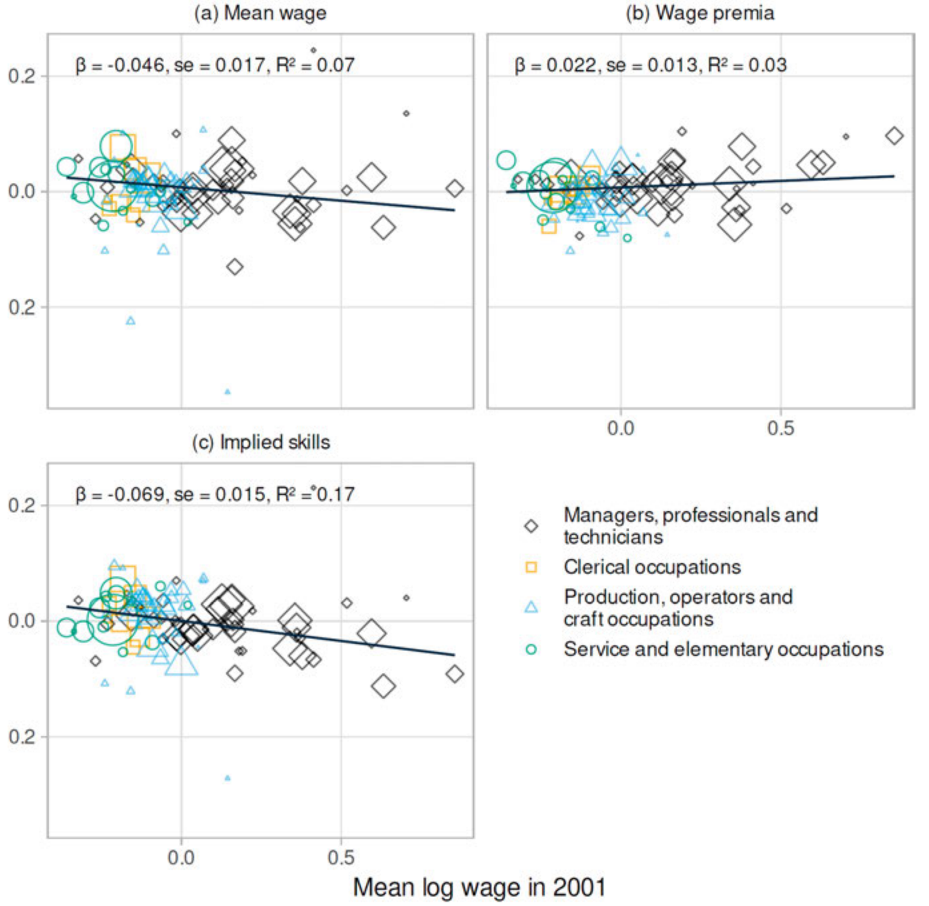

Notes: The figure plots the growth mean log wages, cumulative estimated wage premia, as well as the implied change in mean skills, against initial mean log wages. Wage premia are estimated according to our baseline specification equation (8). Each marker represents one of 101 occupations. The size of each marker is determined by the employment share in the first year and the regression line is weighted accordingly. We use original survey weights when calculating occupation size and mean log wage.

Figure A5: *Premia, skills, and employment growth – robustness checks*

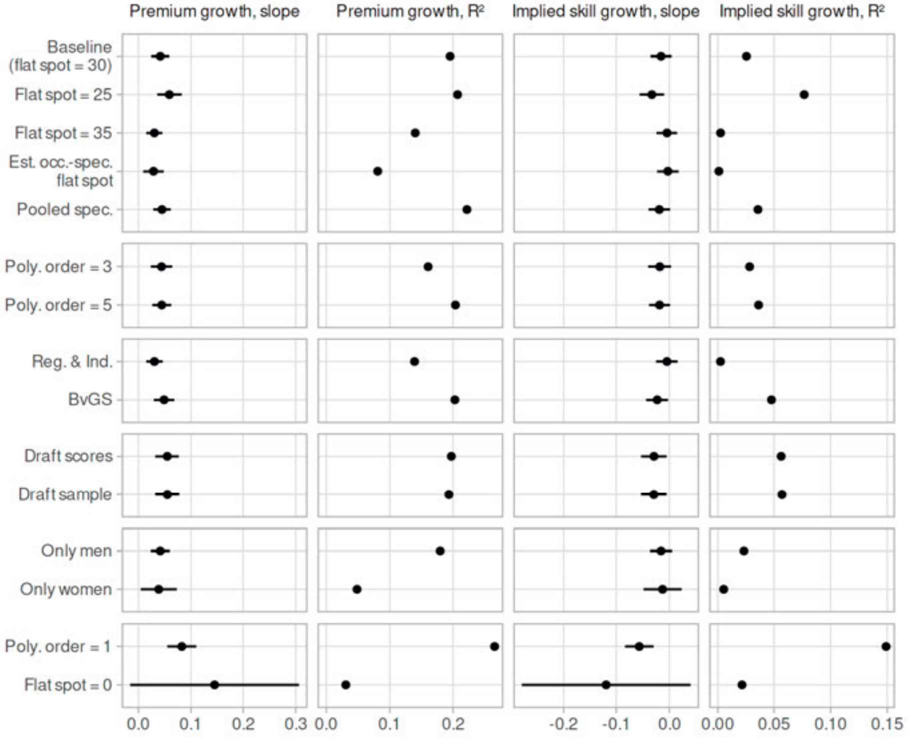

Notes: The table reports the coefficients, standard errors, and coefficients of determination from separately regressing cumulative estimated wage premia and the implied change in mean skills (growth in average wage minus premium growth) against the change in log employment at the occupation level for different sets of premia estimates. See the text for descriptions of how these estimates are produced. The weight assigned to each occupation is determined by the employment share in the first year. We use original survey weights when calculating occupation size and mean log wage.

Figure A6: *Premia, skills, and initial wages – robustness checks*

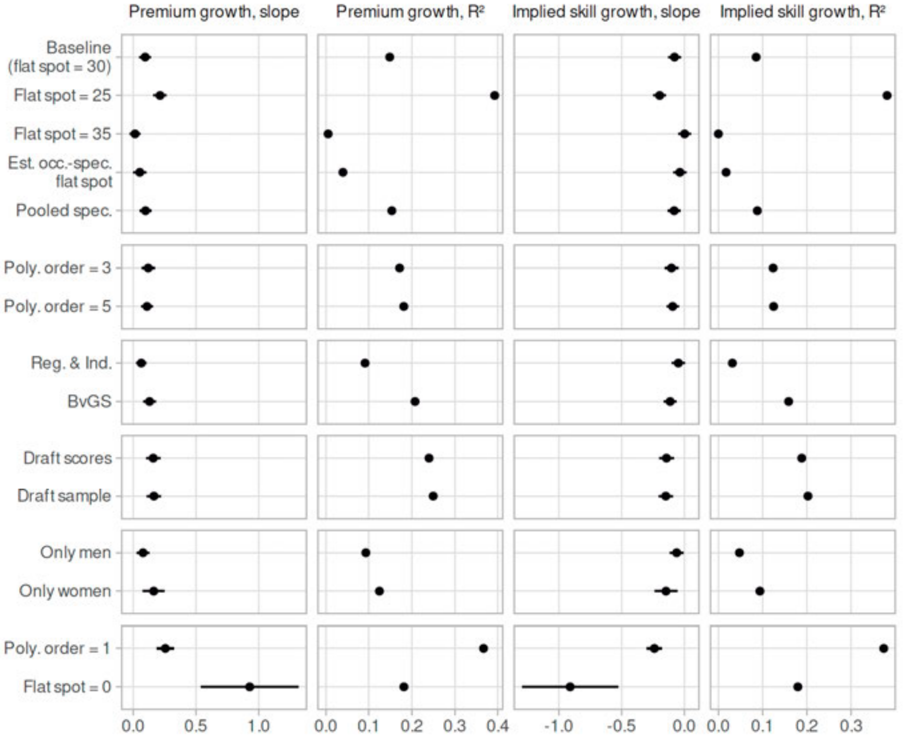

Notes: The figure reports the coefficients, standard errors, and coefficients of determination from separately regressing cumulative estimated wage premia and the implied change in mean skills (growth in average wage minus premium growth) against initial mean log wage at the occupation level for different sets of premia estimates. See also the notes to Figure A5.

Figure A7: *Premia, skills, and schooling – robustness checks*

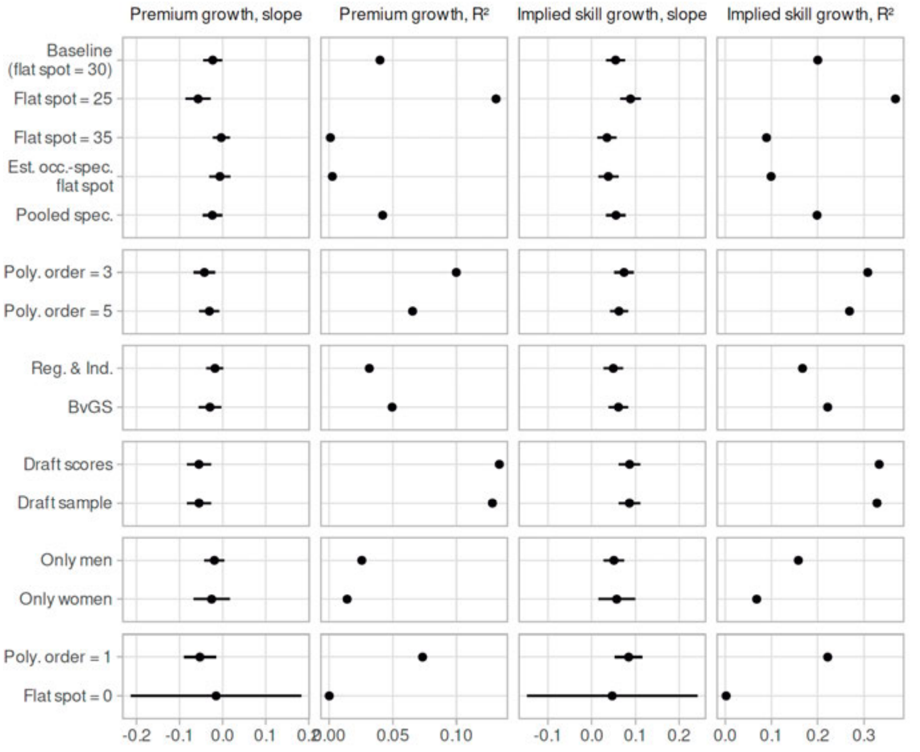

Notes: The figure reports the coefficients, standard errors, and coefficients of determination from separately regressing cumulative estimated wage premia and the implied change in mean skills (growth in average wage minus premium growth) against growth in average years of schooling at the occupation level for different sets of premia estimates. See also the note to Figure A5.

Table A1: Decomposition results—sub-periods

	(1) Baseline	(2) Common flat spot	(3)	(4) Occ.-spec.
		25	35	flat spot
	Panel A: 1996–2013			
Total				
$\Delta \text{Var}(w_{ik})$	2.57			
Between				
$\Delta \text{Var}(w_k)$	1.31			
$\Delta \text{Var}_0(w_k)$	.39			
Components				
$\text{Var}_0(\Delta\pi_k) + 2\text{Cov}_0(w_{k0}, \Delta\pi_k)$	.94	2.03	.29	.66
$\text{Var}_0(\Delta\pi_k)$	.23	.43	.17	.27
$2\text{Cov}_0(w_{k0}, \Delta\pi_k)$	.71	1.59	.12	.4
$\text{Var}_0(\Delta y_k)$	.26	.37	.24	.25
$2\text{Cov}_0(w_{k0}, \Delta y_k)$	-.57	-1.45	.02	-.26
$2\text{Cov}_0(\Delta\pi_k, \Delta y_k)$	-.24	-.55	-.16	-.27
	Panel B: 2001–2013			
Total				
$\Delta \text{Var}(w_{ik})$	.34			
Between				
$\Delta \text{Var}(w_k)$	.07			
$\Delta \text{Var}_0(w_k)$	-.31			
Components				
$\text{Var}_0(\Delta\pi_k) + 2\text{Cov}_0(w_{k0}, \Delta\pi_k)$	.31	.93	-.08	.15
$\text{Var}_0(\Delta\pi_k)$	.09	.13	.08	.1
$2\text{Cov}_0(w_{k0}, \Delta\pi_k)$	.22	.79	-.16	.05
$\text{Var}_0(\Delta y_k)$	.14	.18	.12	.15
$2\text{Cov}_0(w_{k0}, \Delta y_k)$	-.68	-1.25	-.3	-.51
$2\text{Cov}_0(\Delta\pi_k, \Delta y_k)$	-.08	-.17	-.06	-.1

Notes: The table reports results from a decomposition of the change in the between-occupation variance in wages for different flat spot levels and periods. See equation (6) for the formal statement of the decomposition. Column (1) uses a common flat spot for all occupations, at 30 years of potential experience, when estimating growth in wage premia. Columns (2)–(4) vary this common flat spot as indicated. Column (5) estimates a flat spot for each occupation using the procedure described in Appendix A. All figures have been multiplied by 100 for readability.

Table A2: Decomposition results—further specifications

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	
	Baseline	Flat spot = 0	Poly. order			BvGS	Reg. &	Pooled	Draft	Draft	Men	Women
			1	3	5		ind.	spec.	sample	scores		
						Panel A: 1996–2013						
Total												
$\Delta\mathrm{Var}(w_{ik})$	2.57											
Between												
$\Delta\mathrm{Var}(w_k)$	1.31											
$\Delta\mathrm{Var}_0(w_k)$	.39											
Components												
$\mathrm{Var}_0(\Delta\pi_k) + 2\mathrm{Cov}_0(w_{k0}, \Delta\pi_k)$	.94	24.49	2.57	1.21	1.07	1.28	.66	.97	1.64	1.59	.84	2.02
$\mathrm{Var}_0(\Delta\pi_k)$	.23	17.6	.67	.31	.25	.31	.17	.24	.41	.4	.25	.8
$2\mathrm{Cov}_0(w_{k0}, \Delta\pi_k)$	.71	6.89	1.9	.89	.82	.97	.49	.73	1.23	1.19	.59	1.22
$\mathrm{Var}_0(\Delta y_k)$	.26	17.09	.56	.31	.25	.29	.25	.27	.4	.39	.28	.83
$2\mathrm{Cov}_0(w_{k0}, \Delta y_k)$	-.57	-6.75	-1.77	-.75	-.68	-.83	-.35	-.59	-1.09	-1.05	-.45	-1.08
$2\mathrm{Cov}_0(\Delta\pi_k, \Delta y_k)$	-.24	-34.44	-.98	-.37	-.25	-.35	-.18	-.25	-.55	-.54	-.28	-1.38
						Panel B: 2001–2013						
Total												
$\Delta\mathrm{Var}(w_{ik})$	.34											
Between												
$\Delta\mathrm{Var}(w_k)$	.07											
$\Delta\mathrm{Var}_0(w_k)$	-.31											
Components												
$\mathrm{Var}_0(\Delta\pi_k) + 2\mathrm{Cov}_0(w_{k0}, \Delta\pi_k)$	.31	14.5	1.22	.45	.41	.58	.17	.1	.51	.53	.26	.84
$\mathrm{Var}_0(\Delta\pi_k)$	.09	9.31	.22	.12	.09	.13	.07	.11	.09	.08	.11	.29
$2\mathrm{Cov}_0(w_{k0}, \Delta\pi_k)$	.22	5.19	1	.33	.32	.45	.1	-.01	.41	.45	.15	.55
$\mathrm{Var}_0(\Delta y_k)$	.14	9.69	.29	.16	.13	.16	.12	.14	.12	.13	.15	.35
$2\mathrm{Cov}_0(w_{k0}, \Delta y_k)$	-.68	-5.65	-1.45	-.79	-.78	-.91	-.56	-.45	-.87	-.91	-.61	-1.01
$2\mathrm{Cov}_0(\Delta\pi_k, \Delta y_k)$	-.08	-18.85	-.36	-.13	-.08	-.14	-.04	-.1	-.07	-.06	-.11	-.5

Notes: The table reports results from a decomposition of the change in the between-occupation variance in wages for different specifications and periods. See equation (6) for the formal statement of the decomposition and the text for details on the different specifications. All figures have been multiplied by 100 for readability.

Economic Studies

1987:1 Haraldson, Marty. To Care and To Cure. A linear programming approach to national health planning in developing countries. 98 pp.

1989:1 Chryssanthou, Nikos. The Portfolio Demand for the ECU. A Transaction Cost Approach. 42 pp.

1989:2 Hansson, Bengt. Construction of Swedish Capital Stocks, 1963-87. An Application of the Hulten-Wykoff Studies. 37 pp.

1989:3 Choe, Byung-Tae. Some Notes on Utility Functions Demand and Aggregation. 39 pp.

1989:4 Skedinger, Per. Studies of Wage and Employment Determination in the Swedish Wood Industry. 89 pp.

1990:1 Gustafson, Claes-Håkan. Inventory Investment in Manufacturing Firms. Theory and Evidence. 98 pp.

1990:2 Bantekas, Apostolos. The Demand for Male and Female Workers in Swedish Manufacturing. 56 pp.

1991:1 Lundholm, Michael. Compulsory Social Insurance. A Critical Review. 109 pp.

1992:1 Sundberg, Gun. The Demand for Health and Medical Care in Sweden. 58 pp.

1992:2 Gustavsson, Thomas. No Arbitrage Pricing and the Term Structure of Interest Rates. 47 pp.

1992:3 Elvander, Nils. Labour Market Relations in Sweden and Great Britain. A Comparative Study of Local Wage Formation in the Private Sector during the 1980s. 43 pp.

12 Dillén, Mats. Studies in Optimal Taxation, Stabilization, and Imperfect Competition. 1993. 143 pp.

13 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking and Finance. 1993. 303 pp.

14 Mellander, Erik. Measuring Productivity and Inefficiency Without Quantitative Output Data. 1993. 140 pp.

15 Ackum Agell. Susanne. Essays on Work and Pay. 1993. 116 pp.

16 Eriksson, Claes. Essays on Growth and Distribution. 1994. 129 pp.

17 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking and Finance. 2nd version, 1994. 313 pp.

18 Apel, Mikael. Essays on Taxation and Economic Behavior. 1994. 144 pp.

19 Dillén, Hans. Asset Prices in Open Monetary Economies. A Contingent Claims Approach. 1994. 100 pp.

20 Jansson, Per. Essays on Empirical Macroeconomics. 1994. 146 pp.

21 Banks, Ferdinand E.. A Modern Introduction to International Money, Banking, and Finance. 3rd version, 1995. 313 pp.

22 Dufwenberg, Martin. On Rationality and Belief Formation in Games. 1995. 93 pp.

23 Lindén, Johan. Job Search and Wage Bargaining. 1995. 127 pp.

24 Shahnazarian, Hovick. Three Essays on Corporate Taxation. 1996. 112 pp.

25 Svensson, Roger. Foreign Activities of Swedish Multinational Corporations. 1996. 166 pp.

26 Sundberg, Gun. Essays on Health Economics. 1996. 174 pp.

27 Sacklén, Hans. Essays on Empirical Models of Labor Supply. 1996. 168 pp.

28 Fredriksson, Peter. Education, Migration and Active Labor Market Policy. 1997. 106 pp.

29 Ekman, Erik. Household and Corporate Behaviour under Uncertainty. 1997. 160 pp.

30 Stoltz, Bo. Essays on Portfolio Behavior and Asset Pricing. 1997. 122 pp.

31 Dahlberg, Matz. Essays on Estimation Methods and Local Public Economics. 1997. 179 pp.

32 Kolm, Ann-Sofie. Taxation, Wage Formation, Unemployment and Welfare. 1997. 162 pp.

33 Boije, Robert. Capitalisation, Efficiency and the Demand for Local Public Services. 1997. 148 pp.

34 Hort, Katinka. On Price Formation and Quantity Adjustment in Swedish Housing Markets. 1997. 185 pp.

35 Lindström, Thomas. Studies in Empirical Macroeconomics. 1998. 113 pp.

36 Hemström, Maria. Salary Determination in Professional Labour Markets. 1998. 127 pp.

37 Forsling, Gunnar. Utilization of Tax Allowances and Corporate Borrowing. 1998. 96 pp.

38 Nydahl, Stefan. Essays on Stock Prices and Exchange Rates. 1998. 133 pp.

39 Bergström, Pål. Essays on Labour Economics and Econometrics. 1998. 163 pp.

40 Heiborn, Marie. Essays on Demographic Factors and Housing Markets. 1998. 138 pp.

41 Åsberg, Per. Four Essays in Housing Economics. 1998. 166 pp.

42 Hokkanen, Jyry. Interpreting Budget Deficits and Productivity Fluctuations. 1998. 146 pp.

43 Lunander, Anders. Bids and Values. 1999. 127 pp.

44 Eklöf, Matias. Studies in Empirical Microeconomics. 1999. 213 pp.

45 Johansson, Eva. Essays on Local Public Finance and Intergovernmental Grants. 1999. 156 pp.

46 Lundin, Douglas. Studies in Empirical Public Economics. 1999. 97 pp.

47 Hansen, Sten. Essays on Finance, Taxation and Corporate Investment. 1999. 140 pp.

48 Widmalm, Frida. Studies in Growth and Household Allocation. 2000. 100 pp.

49 Arslanogullari, Sebastian. Household Adjustment to Unemployment. 2000. 153 pp.

50 Lindberg, Sara. Studies in Credit Constraints and Economic Behavior. 2000. 135 pp.

51 Nordblom, Katarina. Essays on Fiscal Policy, Growth, and the Importance of Family Altruism. 2000. 105 pp.

52 Andersson, Björn. Growth, Saving, and Demography. 2000. 99 pp.

53 Åslund, Olof. Health, Immigration, and Settlement Policies. 2000. 224 pp.

54 Bali Swain, Ranjula. Demand, Segmentation and Rationing in the Rural Credit Markets of Puri. 2001. 160 pp.

55 Löfqvist, Richard. Tax Avoidance, Dividend Signaling and Shareholder Taxation in an Open Economy. 2001. 145 pp.

56 Vejsiu, Altin. Essays on Labor Market Dynamics. 2001. 209 pp.

57 Zetterström, Erik. Residential Mobility and Tenure Choice in the Swedish Housing Market. 2001. 125 pp.

58 Grahn, Sofia. Topics in Cooperative Game Theory. 2001. 106 pp.

59 Laséen, Stefan. Macroeconomic Fluctuations and Microeconomic Adjustments. Wages, Capital, and Labor Market Policy. 2001. 142 pp.

60 Arnek, Magnus. Empirical Essays on Procurement and Regulation. 2002. 155 pp.

61 Jordahl, Henrik. Essays on Voting Behavior, Labor Market Policy, and Taxation. 2002. 172 pp.

62 Lindhe, Tobias. Corporate Tax Integration and the Cost of Capital. 2002. 102 pp.

63 Hallberg, Daniel. Essays on Household Behavior and Time-Use. 2002. 170 pp.

64 Larsson, Laura. Evaluating Social Programs: Active Labor Market Policies and Social Insurance. 2002. 126 pp.

65 Bergvall, Anders. Essays on Exchange Rates and Macroeconomic Stability. 2002. 122 pp.

66 Nordström Skans, Oskar. Labour Market Effects of Working Time Reductions and Demographic Changes. 2002. 118 pp.

67 Jansson, Joakim. Empirical Studies in Corporate Finance, Taxation and Investment. 2002. 132 pp.

68 Carlsson, Mikael. Macroeconomic Fluctuations and Firm Dynamics: Technology, Production and Capital Formation. 2002. 149 pp.

69 Eriksson, Stefan. The Persistence of Unemployment: Does Competition between Employed and Unemployed Job Applicants Matter? 2002. 154 pp.

70 Huitfeldt, Henrik. Labour Market Behaviour in a Transition Economy: The Czech Experience. 2003. 110 pp.

71 Johnsson, Richard. Transport Tax Policy Simulations and Satellite Accounting within a CGE Framework. 2003. 84 pp.

72 Öberg, Ann. Essays on Capital Income Taxation in the Corporate and Housing Sectors. 2003. 183 pp.

73 Andersson, Fredrik. Causes and Labor Market Consequences of Producer Heterogeneity. 2003. 197 pp.

74 Engström, Per. Optimal Taxation in Search Equilibrium. 2003. 127 pp.

75 Lundin, Magnus. The Dynamic Behavior of Prices and Investment: Financial Constraints and Customer Markets. 2003. 125 pp.

76 Ekström, Erika. Essays on Inequality and Education. 2003. 166 pp.

77 Barot, Bharat. Empirical Studies in Consumption, House Prices and the Accuracy of European Growth and Inflation Forecasts. 2003. 137 pp.

78 Österholm, Pär. Time Series and Macroeconomics: Studies in Demography and Monetary Policy. 2004. 116 pp.

79 Bruér, Mattias. Empirical Studies in Demography and Macroeconomics. 2004. 113 pp.

80 Gustavsson, Magnus. Empirical Essays on Earnings Inequality. 2004. 154 pp.

81 Toll, Stefan. Studies in Mortgage Pricing and Finance Theory. 2004. 100 pp.

82 Hesselius, Patrik. Sickness Absence and Labour Market Outcomes. 2004. 109 pp.

83 Häkkinen, Iida. Essays on School Resources, Academic Achievement and Student Employment. 2004. 123 pp.

84 Armelius, Hanna. Distributional Side Effects of Tax Policies: An Analysis of Tax Avoidance and Congestion Tolls. 2004. 96 pp.

85 Ahlin, Åsa. Compulsory Schooling in a Decentralized Setting: Studies of the Swedish Case. 2004. 148 pp.

86 Heldt, Tobias. Sustainable Nature Tourism and the Nature of Tourists' Cooperative Behavior: Recreation Conflicts, Conditional Cooperation and the Public Good Problem. 2005. 148 pp.

87 Holmberg, Pär. Modelling Bidding Behaviour in Electricity Auctions: Supply Function Equilibria with Uncertain Demand and Capacity Constraints. 2005. 43 pp.

88 Welz, Peter. Quantitative new Keynesian macroeconomics and monetary policy 2005. 128 pp.

89 Ågren, Hanna. Essays on Political Representation, Electoral Accountability and Strategic Interactions. 2005. 147 pp.

90 Budh, Erika. Essays on environmental economics. 2005. 115 pp.

91 Chen, Jie. Empirical Essays on Housing Allowances, Housing Wealth and Aggregate Consumption. 2005. 192 pp.

92 Angelov, Nikolay. Essays on Unit-Root Testing and on Discrete-Response Modelling of Firm Mergers. 2006. 127 pp.

93 Savvidou, Eleni. Technology, Human Capital and Labor Demand. 2006. 151 pp.

94 Lindvall, Lars. Public Expenditures and Youth Crime. 2006. 112 pp.

95 Söderström, Martin. Evaluating Institutional Changes in Education and Wage Policy. 2006. 131 pp.

96 Lagerström, Jonas. Discrimination, Sickness Absence, and Labor Market Policy. 2006. 105 pp.

97 Johansson, Kerstin. Empirical essays on labor-force participation, matching, and trade. 2006. 168 pp.

98 Ågren, Martin. Essays on Prospect Theory and the Statistical Modeling of Financial Returns. 2006. 105 pp.

99 Nahum, Ruth-Aïda. Studies on the Determinants and Effects of Health, Inequality and Labour Supply: Micro and Macro Evidence. 2006. 153 pp.

100 Žamac, Jovan. Education, Pensions, and Demography. 2007. 105 pp.

101 Post, Erik. Macroeconomic Uncertainty and Exchange Rate Policy. 2007. 129 pp.

102 Nordberg, Mikael. Allies Yet Rivals: Input Joint Ventures and Their Competitive Effects. 2007. 122 pp.

103 Johansson, Fredrik. Essays on Measurement Error and Nonresponse. 2007. 130 pp.

104 Haraldsson, Mattias. Essays on Transport Economics. 2007. 104 pp.

105 Edmark, Karin. Strategic Interactions among Swedish Local Governments. 2007. 141 pp.

106 Oreland, Carl. Family Control in Swedish Public Companies. Implications for Firm Performance, Dividends and CEO Cash Compensation. 2007. 121 pp.

107 Andersson, Christian. Teachers and Student Outcomes: Evidence using Swedish Data. 2007. 154 pp.

108 Kjellberg, David. Expectations, Uncertainty, and Monetary Policy. 2007. 132 pp.

109 Nykvist, Jenny. Self-employment Entry and Survival - Evidence from Sweden. 2008. 94 pp.

110 Selin, Håkan. Four Empirical Essays on Responses to Income Taxation. 2008. 133 pp.

111 Lindahl, Erica. Empirical studies of public policies within the primary school and the sickness insurance. 2008. 143 pp.

112 Liang, Che-Yuan. Essays in Political Economics and Public Finance. 2008. 125 pp.

113 Elinder, Mikael. Essays on Economic Voting, Cognitive Dissonance, and Trust. 2008. 120 pp.

114 Grönqvist, Hans. Essays in Labor and Demographic Economics. 2009. 120 pp.

115 Bengtsson, Niklas. Essays in Development and Labor Economics. 2009. 93 pp.

116 Vikström, Johan. Incentives and Norms in Social Insurance: Applications, Identification and Inference. 2009. 205 pp.

117 Liu, Qian. Essays on Labor Economics: Education, Employment, and Gender. 2009. 133 pp.

118 Glans, Erik. Pension reforms and retirement behaviour. 2009. 126 pp.

119 Douhan, Robin. Development, Education and Entrepreneurship. 2009.

120 Nilsson, Peter. Essays on Social Interactions and the Long-term Effects of Early-life Conditions. 2009. 180 pp.

121 Johansson, Elly-Ann. Essays on schooling, gender, and parental leave. 2010. 131 pp.

122 Hall, Caroline. Empirical Essays on Education and Social Insurance Policies. 2010. 147 pp.

123 Enström-Öst, Cecilia. Housing policy and family formation. 2010. 98 pp.

124 Winstrand, Jakob. Essays on Valuation of Environmental Attributes. 2010. 96 pp.

125 Söderberg, Johan. Price Setting, Inflation Dynamics, and Monetary Policy. 2010. 102 pp.

126 Rickne, Johanna. Essays in Development, Institutions and Gender. 2011. 138 pp.

127 Hensvik, Lena. The effects of markets, managers and peers on worker outcomes. 2011. 179 pp.

128 Lundqvist, Heléne. Empirical Essays in Political and Public. 2011. 157 pp.

129 Bastani, Spencer. Essays on the Economics of Income Taxation. 2012. 257 pp.

130 Corbo, Vesna. Monetary Policy, Trade Dynamics, and Labor Markets in Open Economies. 2012. 262 pp.

131 Nordin, Mattias. Information, Voting Behavior and Electoral Accountability. 2012. 187 pp.

132 Vikman, Ulrika. Benefits or Work? Social Programs and Labor Supply. 2013. 161 pp.

133 Ek, Susanne. Essays on unemployment insurance design. 2013. 136 pp.

134 Österholm, Göran. Essays on Managerial Compensation. 2013. 143 pp.

135 Adermon, Adrian. Essays on the transmission of human capital and the impact of technological change. 2013. 138 pp.

136 Kolsrud, Jonas. Insuring Against Unemployment 2013. 140 pp.

137 Hanspers, Kajsa. Essays on Welfare Dependency and the Privatization of Welfare Services. 2013. 208 pp.

138 Persson, Anna. Activation Programs, Benefit Take-Up, and Labor Market Attachment. 2013. 164 pp.

139 Engdahl, Mattias. International Mobility and the Labor Market. 2013. 216 pp.

140 Krzysztof Karbownik. Essays in education and family economics. 2013. 182 pp.

141 Oscar Erixson. Economic Decisions and Social Norms in Life and Death Situations. 2013. 183 pp.

142 Pia Fromlet. Essays on Inflation Targeting and Export Price Dynamics. 2013. 145 pp.

143 Daniel Avdic. Microeconometric Analyses of Individual Behavior in Public Welfare Systems. Applications in Health and Education Economics. 2014. 176 pp.

144 Arizo Karimi. Impacts of Policies, Peers and Parenthood on Labor Market Outcomes. 2014. 221 pp.

145 Karolina Stadin. Employment Dynamics. 2014. 134 pp.

146 Haishan Yu. Essays on Environmental and Energy Economics. 132 pp.

147 Martin Nilsson. Essays on Health Shocks and Social Insurance. 139 pp.

148 Tove Eliasson. Empirical Essays on Wage Setting and Immigrant Labor Market Opportunities. 2014. 144 pp.

149 Erik Spector. Financial Frictions and Firm Dynamics. 2014. 129 pp.

150 Michihito Ando. Essays on the Evaluation of Public Policies. 2015. 193 pp.

151 Selva Bahar Baziki. Firms, International Competition, and the Labor Market. 2015. 183 pp.

152 Fredrik Sävje. What would have happened? Four essays investigating causality. 2015. 229 pp.

153 Ina Blind. Essays on Urban Economics. 2015. 197 pp.

154 Jonas Poulsen. Essays on Development and Politics in Sub-Saharan Africa. 2015. 240 pp.

155 Lovisa Persson. Essays on Politics, Fiscal Institutions, and Public Finance. 2015. 137 pp.

156 Gabriella Chirico Willstedt. Demand, Competition and Redistribution in Swedish Dental Care. 2015. 119 pp.

157 Yuwei Zhao de Gosson de Varennes. Benefit Design, Retirement Decisions and Welfare Within and Across Generations in Defined Contribution Pension Schemes. 2016. 148 pp.

158 Johannes Hagen. Essays on Pensions, Retirement and Tax Evasion. 2016. 195 pp.

159 Rachatar Nilavongse. Housing, Banking and the Macro Economy. 2016. 156 pp.

160 Linna Martén. Essays on Politics, Law, and Economics. 2016. 150 pp.

161 Olof Rosenqvist. Essays on Determinants of Individual Performance and Labor Market Outcomes. 2016. 151 pp.

162 Linuz Aggeborn. Essays on Politics and Health Economics. 2016. 203 pp.

163 Glenn Mickelsson. DSGE Model Estimation and Labor Market Dynamics. 2016. 166 pp.

164 Sebastian Axbard. Crime, Corruption and Development. 2016. 150 pp.

165 Mattias Öhman. Essays on Cognitive Development and Medical Care. 2016. 181 pp.

166 Jon Frank. Essays on Corporate Finance and Asset Pricing. 2017. 160 pp.

167 Ylva Moberg. Gender, Incentives, and the Division of Labor. 2017. 220 pp.

168 Sebastian Escobar. Essays on inheritance, small businesses and energy consumption. 2017. 194 pp.

169 Evelina Björkegren. Family, Neighborhoods, and Health. 2017. 226 pp.

170 Jenny Jans. Causes and Consequences of Early-life Conditions. Alcohol, Pollution and Parental Leave Policies. 2017. 209 pp.

171 Josefine Andersson. Insurances against job loss and disability. Private and public interventions and their effects on job search and labor supply. 2017. 175 pp.

172 Jacob Lundberg. Essays on Income Taxation and Wealth Inequality. 2017. 173 pp.

173 Anna Norén. Caring, Sharing, and Childbearing. Essays on Labor Supply, Infant Health, and Family Policies. 2017. 206 pp.

174 Irina Andone. Exchange Rates, Exports, Inflation, and International Monetary Cooperation. 2018. 174 pp.

175 Henrik Andersson. Immigration and the Neighborhood. Essays on the Causes and Consequences of International Migration. 2018. 181 pp.

176 Aino-Maija Aalto. Incentives and Inequalities in Family and Working Life. 2018. 131 pp.

177 Gunnar Brandén. Understanding Intergenerational Mobility. Inequality, Student Aid and Nature-Nurture Interactions. 2018. 125 pp.

178 Mohammad H. Sepahvand. Essays on Risk Attitudes in Sub-Saharan Africa. 2019. 215 pp.

179 Mathias von Buxhoeveden. Partial and General Equilibrium Effects of Unemployment Insurance. Identification, Estimation and Inference. 2019. 89 pp.

180 Stefano Lombardi. Essays on Event History Analysis and the Effects of Social Programs on Individuals and Firms. 2019. 150 pp.

181 Arnaldur Stefansson. Essays in Public Finance and Behavioral Economics. 2019. 191 pp.

182 Cristina Bratu. Immigration: Policies, Mobility and Integration. 2019. 173 pp.

183 Tamás Vasi. Banks, Shocks and Monetary Policy. 2020. 148 pp.

184 Jonas Cederlöf. Job Loss: Consequences and Labor Market Policy. 2020. 213 pp.

185 Dmytro Stoyko. Expectations, Financial Markets and Monetary Policy. 2020. 153 pp.

186 Paula Roth. Essays on Inequality, Insolvency and Innovation. 2020. 191 pp.

187 Fredrik Hansson. Consequences of Poor Housing, Essays on Urban and Health Economics. 2020. 143 pp.

188 Maria Olsson. Essays on Macroeconomics: Wage Rigidity and Aggregate Fluctuations. 2020. 130 pp.

189 Dagmar Müller. Social Networks and the School-to-Work Transition. 2020. 146 pp.

190 Maria Sandström. Essays on Savings and Intangible Capital. 2020. 129 pp.

191. Anna Thoresson. Wages and Their Impact on Individuals, Households and Firms. 2020. 220 pp.

192. Jonas Klarin. Empirical Essays in Public and Political Economics. 2020. 129 pp.

193. André Reslow. Electoral Incentives and Information Content in Macroeconomic Forecasts. 2021. 184 pp.

194. Davide Cipullo. Political Careers, Government Stability, and Electoral Cycles. 2021. 308 pp.

195. Olle Hammar. The Mystery of Inequality: Essays on Culture, Development, and Distributions. 2021. 210 pp.

196. J. Lucas Tilley. Inputs and Incentives in Education. 2021. 184 pp.

197. Sebastian Jävervall. Corruption, Distortions and Development. 2021. 215 pp.

198. Vivika Halapuu. Upper Secondary Education: Access, Choices and Graduation. 2021. 141 pp.

199. Charlotte Paulie. Essays on the Distribution of Production, Prices and Wealth. 2021. 141 pp.

200. Kerstin Westergren. Essays on Inflation Expectations, Monetary Policy and Tax Reform. 2021. 124 pp.

201. Melinda Süveg. Finance, Shocks, Competition and Price Setting. 2021. 137 pp.

202. Adrian Poignant. Gold, Coal and Iron. Essays on Industrialization and Economic Development. 2022. 214 pp.

203. Daniel Bougt. A Sequence of Essays on Sequences of Auctions. 2022. 188 pp.

204. Lillit Ottosson. From Welfare to Work. Financial Incentives, Active Labor Market Policies, and Integration Programs. 2022. 219 pp.

205. Yaroslav Yakymovych. Workers and Occupations in a Changing Labour Market. The Heterogeneous Effects of Mass Layoffs and Social Safety Nets. 2022. 212 pp.