Institute for International Economic Studies
Monograph Series No. 118

Stockholm University

Essays on Home Production, Mobility, and Monetary Policy

Mattias Almgren

Doctoral Thesis in Economics at Stockholm University, Sweden 2023

Essays on Home Production, Mobility, and Monetary Policy

Mattias Almgren

Academic dissertation for the Degree of Doctor of Philosophy in Economics at Stockholm University to be publicly defended on Wednesday 24 May 2023 at 13.00 in Nordenskiöldsalen, Geovetenskapens hus, Svante Arrhenius väg 12.

Abstract

The Allocation of Expenditures and Time over Time

In the year 2018, the average high skilled single male worked more than 35 hours per week in the market and allocated more than 70 percent of expenditures to services. In the same year, the average low skilled female worked 22 hours per week in the market and allocated around 65 percent of expenditures to services. What might explain these differences at that point in time, and how they have developed over time? In this paper, I compile and analyze data related to time-use and the allocation of expenditures. The dataset that I compile is customized to fit the needs of my companion paper, *Home Production, Expenditures, and Time Use*, in which I propose and test a theory that includes home production.

Home Production, Expenditures, and Time Use

I propose a unified analysis of expenditures and time use, via a model in which households can produce services at home. I show that the model, which uses stable homothetic preferences and standard functional forms for home production, can match data for the U.S. about expenditures and time-use, both in the cross-section and the developments over time. For women, changes in social norms were important. Absent changes in social norms, the developments would have been vastly different, both in terms of how they allocated their time and in terms of how expenditures were allocated.

It Runs in the Family: Occupational Choice and the Allocation of Talent

Children frequently grow up to work in the same jobs as their parents. Using unique data on worker skills and personality traits, and administrative data on the labor market outcomes of Swedish men, we study how skills and parental background influence occupational choice, intergenerational mobility, and the allocation of talent in the economy. First, we document that sons are disproportionately more likely to follow into the same occupation as their fathers, across all skills and earnings levels. Second, we estimate a general equilibrium Roy model with costly occupational choice and heterogeneous entry barriers depending on parental background. We find that these entry barriers lead to misallocation: Equalizing entry costs across workers leads occupational following to fall by half. This leads to increased intergenerational mobility, with the largest income gains among sons of fathers in the bottom income decile. Third, exploiting structural change in employment in fathers' occupations, difference-in-differences estimates imply that occupational following leads to reduced earnings, concentrated among sons of low-income fathers and those whose skills are misaligned with those of incumbents in their father's occupation. Our findings suggest that equalizing career opportunities bring equity gains.

Monetary Policy and Liquidity Constraints: Evidence from the Euro Area

We quantify the relationship between the response of output to monetary policy shocks and the share of liquidity constrained households. We do so in the context of the euro area, using a Local Projections Instrumental Variables estimation. We construct an instrument for changes in interest rates from changes in overnight indexed swap rates in a narrow time window around ECB announcements. Monetary policy shocks have heterogeneous effects on output across countries. Using micro data, we show that the elasticity of output to monetary policy shocks is larger in countries that have a larger fraction of households that are liquidity constrained.

Keywords: *Macroeconomics, monetary economics, monetary policy, misallocation, intergenerational mobility, structural change, home production, labor supply.*

Stockholm 2023
http://urn.kb.se/resolve?urn=urn:nbn:se:su:diva-216083

ISBN 978-91-8014-270-0
ISBN 978-91-8014-271-7
ISSN 0346-6892

Stockholm University

Department of Economics

Stockholm University, 106 91 Stockholm

ESSAYS ON HOME PRODUCTION, MOBILITY, AND MONETARY POLICY

Mattias Almgren

Stockholm University

Essays on Home Production, Mobility, and Monetary Policy

Mattias Almgren

ISBN print 978-91-8014-270-0
ISBN PDF 978-91-8014-271-7
ISSN 0346-6892

Cover photo: Picture from rawpixel.com, picture ID 5945862
Back-cover photo by Hanna Weitz, Stockholm University
Printed in Sweden by Universitetsservice US-AB, Stockholm 2023

To my family.

Doctoral Dissertation
Department of Economics
Stockholm University

Abstracts

The Allocation of Expenditures and Time over Time
In the year 2018, the average high skilled single male worked more than 35 hours per week in the market and allocated more than 70 percent of expenditures to market services. In the same year, the average low skilled female worked 22 hours per week in the market and allocated around 65 percent of expenditures to market services. What might explain these differences at that point in time, and how they have developed over time? In this paper, I compile and analyze data related to time-use and the allocation of expenditures. The dataset that I compile is customized to fit the needs of my companion paper, *Home Production, Expenditures, and Time Use*, in which I propose and test a theory that includes home production.

Home Production, Expenditures, and Time Use
I propose a unified analysis of expenditures and time-use, via a model in which households can produce services at home. I that the model, which uses stable homothetic preferences and standard functional forms for home production, can match data for the U.S. about expenditures and time-use, both in the cross-section and the developments over time. For women, changes in social norms were important. Absent changes in social norms, the developments would have been vastly different, both in terms of how they allocated their time and in terms of how expenditures were allocated.

It Runs in the Family: Occupational Choice and the Allocation of Talent

Children frequently grow up to work in the same jobs as their parents. Using unique data on worker skills and personality traits, and administrative data on the labor market outcomes of Swedish men, we study how skills and parental background influence occupational choice, intergenerational mobility, and the allocation of talent in the economy. First, we document that sons are disproportionately more likely to follow into the same occupation as their fathers, across all skills and earnings levels. Second, we estimate a general equilibrium Roy model with costly occupational choice and heterogeneous entry barriers depending on parental background. We find that these entry barriers lead to misallocation: Equalizing entry costs across workers leads occupational following to fall by half. This leads to increased intergenerational mobility, with the largest income gains among sons of fathers in the bottom income decile. Third, exploiting structural change in employment in fathers' occupations, difference-in-differences estimates imply that occupational following leads to reduced earnings, concentrated among sons of low-income fathers and those whose skills are misaligned with those of incumbents in their father's occupation. Our findings suggest that equalizing career opportunities bring equity gains.

Monetary Policy and Liquidity Constraints: Evidence from the Euro Area

We quantify the relationship between the response of output to monetary policy shocks and the share of liquidity constrained households. We do so in the context of the euro area, using a Local Projections Instrumental Variables estimation. We construct an instrument for changes in interest rates from changes in overnight indexed swap rates in a narrow time window around ECB announcements. Monetary policy shocks have heterogeneous effects on output across countries. Using micro data, we show that the elasticity of output to monetary policy shocks is larger in countries that have a larger fraction of households that are liquidity constrained.

Acknowledgments

I will be forever grateful to my two supervisors Per Krusell and Timo Boppart. Without their encouragement, I would probably not have achieved anything as an innebandy player. Without their encouragement, I would have much more often, when I felt that I had many other things on my plate, skipped innebandy practice. I am certain that this type of encouragement, especially coming from the both of them, helped balance my life in a positive way, and that it has been an important contributing factors towards that I am now able to finally submit this thesis.

Pursuing a PhD degree has been more demanding than I imagined that it would be. For me, the part that has been the most difficult has been the feeling that, even though breakthroughs do occur and eventually might take a project over the finish line, setbacks and problems take up more time and focus. Per and Timo have been very important for me, as they have encouraged me to proceed, and find solutions, when I thought that I had encountered insurmountable obstacles. Sometimes, they made me realize that the problems were not as big as I imagined, and when I was faced with actual problems, they consistently steered me in the right direction.

I am glad for that I have spent my time at Stockholm University and, after my second year, at the IIES, which truly is a special place. I feel very blessed to have gotten the opportunity to interact and work with many talented economists. First, I would like to thank my coauthors, starting with John Kramer, who is my co-author in two of the papers which are included in this dissertation. I think that he has developed from a great economist to an ever greater one, and I think that we have complemented each other very well. One of the papers in this dissertation would not have come as far as it has without the skillful work of Ricardo Lima. I do also want to thank my co-author José-Elías Gallegos for all the hard work that he has put in, and for being a person who focuses on what is positive in a world with economists who often tend to focus on what is not good enough. I have also had the fortune to work with Jósef Sigurdsson, who's conviction and research confidence has been contagious.

My first years in the program would not have been the same without the fantastic group of people that entered it at the same time. I would especially like to thank Christine Alamaa, Iman Dadgar, Roza Khoban, Jon Olofssen, and Malin Tallås Ahlzén. Xuepin Sun also started with us at the same time, and I was blessed with having an office close to hers for many years at the IIES. She amazed me already during the first weeks of the program, and never stopped doing that.

There are many other, both current and former, students at the IIES who have all contributed to my resaerch: Gualtiero Azzalini, Agneta Berge, Max Bres, Saman Darougheh, Mohammadreza Farajpour, Richard Foltyn, Ida Haavi, Karl Harmenberg, Stefan Hinkelmann, Philipp Hochmuth, Markus Karlman, Matilda Kilström, Karin Kinnerud, Kasper Kragh Balke, Markus Kondziella, Carolina Lindholm, Jonna Olsson, Fredrik Paues, Markus Peters, Sreyashi Sen, Claire Thürwächter, Has van Vlokhoven, Magnus Åhl, and Erik Öberg. Thank you! Then there are many colleagues that I have not interacted with as much when it comes to research, but have been important for making my time in the program enjoyable: Iacobo Bianchi, Mattias Folkestad, Evelina Linnros, Francesco Loiacono, Patrizia Massner, Fabian Sinn, Sebestian Tebbe, and Tillman von Carnap-Bornheim.

Then I am grateful for all the advice that I have gotten from faculty members. John Hassler and Mitchell Downey have been sitting in offices close to mine, and their doors have always stood open and they have been willing to discuss whatever ideas that I might have had. And there are many others that have given me advice related to my research that I would like to thank: Tobias Broer, Konrad Burchardi, Jonathan de Quidt, Kieran Larkin, Kurt Mitman, Arash Nekoei, Peter Nilsson, Torsten Persson, David Schönholzer, David Seim, Joshua Weiss, and Yimei Zou.

Many others are responsible for making the IIES a place where you want to be. For this I would like to thank Ingvild Almås, Laia Narro-Sola, Mats Persson, and Tessa Bold. I would like to especially thank Jakob Svensson, director at the IIES, for letting, and making, the IIES be a place where a lot of fun is had. The place would also not be the same without Tove Lundström, and Hanna Weitz. Then I would like to give a big thank you to Ulrika Gålnander, among other things for being of big help during these final stages. Another person who I will be forever grateful for is Christina Lönnblad, who I think is one of the most reliable people I have ever met. And I also owe her a big thank you for all the help that she has given me at these final stages towards completing my dissertation.

I have had colleagues over the years, outside academia, that I would like to take the opportunity to thank. For always have been encouraging me and for making me feel appreciated, I would like to thank Anna Breman, Anna Felländer, Knut Hallberg, Olof Manner, Gill Nacksten-Marcusson, and Håkan Persson. Then I would also like to thank Thomas Eisensee, for all his support and for being so incredibly considerate during these final stages.

I have many friends that I would like to thank. I would especially like to thank Max and Johanna, who have been amazing friends for many years now. Without their love and support, this would not have

been possible. And then I want to say thank you to Ina and Gustav, who always have shown their interest in and support for what I have been trying to do. I am also very happy that I have a-enheten in my life, and would like to thank Anders, André, David, Jacob, 3xJohan, Ludwig, Mathias, Nicklas, Oscar, Viktor, and Zjenja. Around you guys, I always manage to totally forget about work for a while, and for this I am really grateful.

During tough times, I have always felt that I could rely on my family for support. Sometimes they have been people for me to talk with about my struggles, sometimes my family has just been a big group of loving people in which I can just be myself and reload. I especially want to say thank you to my parents, Carin and Michael, for being there, and never showing any doubts about what I can do. I would also like to thank my parents in law, Anna and Staffan, who also have supported me and my family tremendously. Anna has helped us get our everyday life together, and I do not know how we would have managed without her.

During the second year in the PhD program, I had my first daughter, Saga, and then, a couple of years later, my second daughter, Ilse. I am just amazed by how fast time has gone by; at the time of writing this, Ilse and Saga are three and six years old. Without them, my everyday life would not be filled with action, laughter, and meaning. I am so immensely grateful for having them in my life, and thank them from the bottom of my heart for all the love that they have been giving me.

Matilda has been my partner for many years, and my wife since the beginning of the second year in the PhD program. It is difficult to know what to write to Matilda; how do I thank her for all that she has done for me over the years? I love you so incredibly much, and I am looking forward to exploring what the future has in store for us.

Mattias Almgren
Stockholm, Sweden
April 2023

Contents

Introduction

This dissertation contains four self-contained chapters. The first two chapters analyze and try to understand how households in the U.S. have spent their time, and how they have allocated their expenditures on different types of consumption. In the third chapter, the importance of family background for occupational choice and its implications for intergenerational earnings mobility is studied. The fourth chapter estimates differential effects of monetary policy on output in countries in the euro area and investigates potential explanations for it.

The Allocation of Expenditures and Time over Time

This chapter starts from the hypothesis that it is important to consider the relationship between how households allocate their time and expenditures. I especially stress the importance of studying what households do when they are not working in the market; are they enjoying leisure (which standard models with labor supply would assume), or are they working in home production? First, it collects, processes, summarizes, describes, and analyzes economic data from the U.S. about how households allocate their time and expenditures. This serves a purpose on its own, as it allows for a unified analysis of these dimensions. In my second chapter I develop a theory that focuses on how households allocate their time and expenditures, and I emphasize the potential importance of home production. Hence, a second purpose of this chapter is to produce a dataset that can be used in Chapter 2.

A well-known fact about the data is that the expenditure share on services has been increasing for many decades. I document that this increase is present in all types of households that I consider. Moreover, all these types of households increased their consumption of market services relative to nondurable goods over many decades while the relative price of market services increased, which is something that the standard economic theory with stable homothetic preferences cannot account for.

The average hours worked in the working age population has increased since the early 1960s. However, this average is the result of developments that are different in different types of households. First, average hours worked among men have decreased over time, while they have increased considerably among women. While the decrease is widespread across groups of men, the increase in hours worked among women is to a large part explained by women in couples having increased their market labor supply. At the same time, hours worked in home production have increased among men over time, while they have decreased among women, especially among women in couples.

I summarize and compare how wage rates have developed for different types of men and women. On average, women's wage rates have been increasing more than wage rates among men. But looking more closely at the data, and analyzing wage rates jointly with hours worked at home and in the market, makes it clear that movements in wage rates, on their own, are likely not enough to explain the data. E.g., wage rates among low skilled single women have been consistently below the wage rates among low skilled single men. Yet, expenditure shares on services have been strikingly similar between these two groups. Moreover, the relative wage rate of low skilled single women relative to the wage rate of low skilled single men has not changed over time. Despite this, this group of women have increased how much labor they supply to the market, while the opposite is true among low skilled men. In my next chapter, I suggest that it could be important to consider home production and social norms when seeking to understand the allocation of time and expenditures.

Home Production, Expenditures, and Time Use

In Chapter 1, I document that a significant share of individuals' time budgets is allocated to home production. I also document that how much men and women work in home production has changed considerably over time, and the differences between different types of households have been substantial. If households can, to some extent, substitute between market and home produced goods and services, then hours worked in the market, and expenditures on market services and goods, might change if the cost of producing output at home changes. In Chapter 2, I examine whether home production can help explain patterns in the allocation of expenditures and time.

There are several facts about the data that a standard economic model with stable homothetic preferences cannot account for. One is the fact that households increased their consumption of market services relative to nondurable goods over many decades while the relative price of market services increased. A popular, and reasonable, way to be able to account for this fact is to assume that households have non-homothetic preferences. In short, they assume that as income grows there is an income effect that induces households to spend a greater part of their expenditures on services. The purpose of this chapter is to test whether a model with stable homothetic preferences and home production can account for the data, and if so under what conditions.

In my framework households consume services, which can be purchased in the market, or produced at home by households by combining their own time with capital and nondurable goods. E.g., to produce meatballs at home, the households need to combine their time with ingredients (e.g., minced meat and spices) and capital (e.g., stove and frying pan). To finance expenditures on market services, nondurable goods and capital, they must earn labor income, which they get by selling their time to the market.

I focus on low and high skilled single men and women, respectively. An ex-ante plausible story concerning why the consumption of market services increased relative to the consumption of nondurable goods is that households substituted away from home produced services to market services. For men, however, I find this not to have been the case. Rather, I find that men substituted in the other direction – to home services. This was driven by the fact that it became cheaper to produce home services, which was partly driven by decreased prices of nondurable goods and capital, and partly by increased efficiency. Interestingly, I find that men substituted away considerably from nondurable goods in home production, and that this explains why the consumption of market services increased relative to the consumption of nondurable goods.

For the same reasons, women also substituted away from nondurable goods in home production. Unlike men, however, they did substitute from home production to market production. The factor underlying this development is changes is social norms. These social norms are modeled as affecting the relative disutility from working in the market relative to working in home production, and I estimate the changes in social norms to have been substantial. More specifically, I estimate that social norms acted as a barrier for women to enter the labor market and as something that kept them in home production. These changes in social norms made women, over time, substitute from spending time in home production to market work. Other than having direct and important effects on time-use, social norms made women demand more market services and reduced the need for nondurable goods used in home production.

It Runs in the Family: Occupational Choice and the Allocation of Talent

It is a well-known fact that many children end up in the same occucaptions as one of their parents. For example, using Swedish data, John Kramer, Josef Sigurdsson, and I show that sons of medical doctors are about ten times as likely to become doctors themselves, as compared to sons with fathers in other occupations. This probability bias, which we refer to as the occupational mobility bias, is on average six among men in Sweden.

Why is this bias so large? By trying to answer this question, we enter a large literature with a long history. One part of the literature has been documenting the same type of phenomenon, i.e., that children are likely to follow in their parents' footsteps. Another, but not completely separate literature, focuses on to what extent outcomes, specifically in our case related to occupational choices and incomes, are a result of innate abilities, on the one hand, and nurture on the other. That is, do children often gravitate towards the same occupations as their parents because they posses similar abilities, or is it something else?

We find that occupational choice has large implications for the intergenerational correlation in earnings ranks that we observe, as it turns out that this correlation can be fully explained by how the sons sort across occupations. Hence, if we learn more about the drivers behind occupational choice, we can better understand the drivers behind intergenerational earnings mobility.

Using data on actual earnings and measured skills, we do for each individual in our dataset estimate potential earnings across the full set of occupations. We find that the measured skills can explain some of the bias that we observe in terms of occupational following, but that a big part of the bias must be explained by something else. We label these other things as "discounts", but identifying what they reflect is difficult, and it could be anything from unequal access to information and education (e.g., medical doctors and lawyers), experience and knowledge (e.g., farmers, lawyers, and entrepreneurs), networks and connections (e.g., corporate managers and politicians), to family firms (e.g., pharmacists). We estimate these discounts using a structural framework, and then run a counterfactual experiment which entails removing them. As the discounts are removed, occupational following does decrease considerably. While occupational following is 8.6 percent in the data and the baseline, it drops to 3.4 percent in the counterfactual economy. Interestingly, while many individuals make other choices with regards to which occupation to enter, the effect on aggregate output is close to zero. In terms of the intergenerational income rank, it decreases from 0.245 to 0.217, and it is especially sons to fathers in the lowest part of the income distribution that gain the most.

To complement the results from the structural model, we estimate a couple of empirical models. In these, we estimate what happens to earnings for sons whose fathers' occupations are in decline. When the occupations of their fathers are shrinking in size, the probability of following into those occupations falls. The idea is then to estimate how and to what extent this affects earnings. Our results indicate that sons who have fathers in occupations that decline do have higher earnings than what would otherwise have been the case. Moreover, we find that this positive effect comes from (i) sons whose fathers are in the lower half of the income distribution, and (ii) whose skill profiles are less fit for working in the father's occupation. The decrease in the size of the father's occupation could be seen as something that implicitly increases the incentives for the son to pursue another occupation. Interestingly, our results suggest that it is in particular sons who would otherwise be a poor fit in the fathers' occupations that are making other choices, and that this is especially the case among sons of fathers at the lower end of the income distribution.

Monetary Policy and Liquidity Constraints: Evidence from the Euro Area

The main objective for most modern central banks is price stability, which is believed to reduce business cycle fluctuations and improve the conditions for economic growth. When inflation is on the rise, the central bank increases the policy rate, which is intended to reduce inflation partly through decreased demand. But how much does the central bank policy rate affect demand and output? Moreover, do the magnitudes of these effects vary across countries, and if so, why? These are the questions that I, together with my co-authors José-Elías Gallegos, John Kramer, and Ricardo Lima, pose in this paper.

To answer these questions, we make use of the fact that there is a large number of countries that are all directly subject to the same policy rate set by a common central bank – namely countries in the euro area. Focusing on the euro area is ideal for our purposes for two reasons. First, since member countries share a common central bank, we only have to identify one set of monetary policy shocks. This means that any estimated differences in output responses can be attributed to differential effects on output in the different countries, and not to the extent monetary policy shocks were more or less well-identified. Second, for various historical reasons, the countries differ from each other in many dimensions. Naturally, this is a necessary condition for answering the question about *why* output might be more affected in some countries than in others.

We identify so-called monetary policy shocks using high-frequency financial data, and then estimate the effects on output using modern econometric methods. Interestingly, we find that output-responses are substantially different across countries in the euro area. Using data from the Household Finance and Consumption Survey (HFCS), which collects household-level data on households' finances and consumption, we compute various country-level statistics, which we then relate to the estimated output responses.

A number of papers in the theoretical literature have earlier highlighted households with little liquid wealth as potentially being important for the effectiveness of monetary policy. These are households that might be extra sensitive to fluctuations in income and who are more likely to have to adjust their levels of consumption, rather than smoothing consumption by tapping into a liquid savings account. Following some influential contributions to the theoretical literature, one of the statistics that we calculate is the share of Hand-to-Mouth households (HtM) in each respective country, which essentially are households with low levels of liquidity. Among the different statistics that we consider, the HtM shares are the only ones that seem to matter for how much output responds to monetary policy shocks; we find that output responds more to monetary policy shocks in countries in which the HtM shares are larger.

Our results have implications for research in economics. In particular, some earlier theoretical papers have shown that the share of HtM households has ambiguous effects on monetary policy effectiveness, depending on whether the income elasticity with respect to aggregate income is higher among HtM households than among non-HtM households. Our results can guide future theoretical work in this field. The results are also valuable for policy makers. Specifically, we show that central bankers should consider the share of liquidity constrained households when deciding on what is a well-balanced policy rate. Moreover, advisors to and members of the Governing Council at the European Central Bank need to be aware of their policy having heterogeneous effects on output in the different countries in the euro area.

Chapter 1

The Allocation of Expenditures and Time Over Time*

*I am grateful to Timo Boppart and Per Krusell for their advice and support. I would also like to take the opportunity to thank Mitch Downey, Richard Foltyn, Jose-Elias Gallegos, Stefan Hinkelmann, Philipp Hochmuth, John Kramer, Kieran Larkin, Kurt Mitman, Jonna Olsson, David Schoenholtzer, Josef Sigurdsson, Xueping Sun, and Has van Vlokloven.

1.1 Introduction

There are large systematic differences between individuals from different types of households in terms of how they allocate time and expenditures. In the year 2018, the average high skilled single male worked more than 35 hours per week in the market and allocated more than 70 percent of expenditures to market services. In the same year, the average low skilled female worked 22 hours per week in the market and allocated around 65 percent of expenditures to market services. What might explain these differences at that point in time, and how they have developed over time?

For most working age individuals, market work takes up a significant fraction of the weekly time budget. At the firm, hours are combined with other factors of production and produce output. For many decades now, methods have been developed to measure the value added of what is produced, which, at an aggregated level, but also at finer levels, is rightfully of major interest to policy makers. Since labor supplied by individuals to firms can account for a large part of the value added that is generated, it is of great importance to both measure and understand it.

Hours that are not aggregated to work in the market are usually aggregated into one residual category labeled leisure, and much less focus has been allocated towards a more detailed understanding of this residual category and its effect on the economy, and welfare. A large fraction of these hours are spent in home production. In the year 2018, the average low skilled single woman worked 18 hours in home production, i.e., spent just four hours less than time spent on market work. How expenditures are allocated might be connected to what individuals do when they are not working in the market. I take particular interest in the potential importance of home production for understanding both cross-sectional differences and time-series developments in market hours, as well as how individuals allocate expenditures between different types of goods and services.

One purpose of this paper is to customize a coherent dataset that can be used in my companion paper Almgren (2023), in which I develop a theory for understanding facts about time-use and expenditures. As a simple, but potent, motivation for doing that, in this paper I show that some of the predictions from a standard economic framework with stable homothetic preferences do not square with what is observed in the data. I document cross-sectional differences, and developments over time, about time-use and allocations of expenditures. Few of the observations are novel, as a long list of papers have had a similar focus, but few of them consider all these dimensions simultaneously. A comprehensive literature summary is plainly infeasible; I instead make some comparisons when I present and analyze the data.

The companion paper investigates whether a model with stable homothetic preferences and home production potentially can generate results that align with time-use and expenditure data. Specifically, the paper focuses on low and high skilled single men and women, and how time-use, and expenditures on different types of consumption groups, are different in the cross-section, and how they have developed over time. In this paper, I also compile data for couple households. The reason for this is twofold. First, observing and analyzing the empirical regularities related to time-use and expenditures for couples is interesting on its own. Second, a natural extension to Almgren (2023) is to model couples. Taking that model to the data is easier when the data has already been put together.

Based on the setup of the model, time-use is allocated between three activities: (i) market work, (ii) domestic work, and (iii) leisure. Expenditures can be allocated towards three types of consumption: (i) market services, (ii) nondurable goods, and (iii) durable goods (capital). Lastly, I need prices, both for consumption, and also wage rates. All prices are expressed relative to the price of market services, i.e., the price of market services is chosen as the numeraire. As the model in Almgren (2023) is static and does not include initial assets, unearned income, nor savings, I make adjustments to expenditures for each household in each period such that the budget constraint holds.

First, in Section 1.2, I look at how the share of individuals with at least a college degree has developed since the beginning of the 1960s, in the aggregate and separately for men and women in single and couple households, respectively. I do also investigate some patterns about sorting. Whether to couple with someone, and if so, who that would be, is not modeled in Almgren (2023). Nevertheless I look at these patterns, as they highlight a dimension that could be important and could be considered in an extended framework. The rest of Section 1.2 then goes on with describing the data sources used and how the data is put together. Section 1.2.2 looks at labor supply in the market, and then Section 1.2.3 presents details for how to get time series for expenditures at the different consumption groups for the different types of households. In Section 1.2.4, I present how I calculate hourly wage rates, and Section 1.2.5 presents data on hours worked at home. Section 1.3 makes some observations about how labor supply and wage rates have developed relative to each other for men and women. Then in Section 1.4 I highlight some stylized facts that a standard economic framework cannot explain. Section 1.6 concludes the paper.

1.2 Data Construction and Analysis

1.2.1 Household Types

I use data from the Current Population Survey between the years 1962-2018, downloaded from IPUMS. An individual is uniquely identified by the combination of four variables: (i) *year*, (ii) *month*, (iii) *serial*, and (iv) *pernum*. The prior two provide information about in which year and month the survey took place, and *serial* and *pernum* refer to a specific household, and member in that household, respectively.

First, I classify individuals as either (i) "household heads", or (ii) "spouses". An individual without a cohabiting partner is classified as "single household" and naturally as a household head. I define a "couple household" as a household that consists of two co-habiting partners. More specifically, the relationship between the household head and each interviewed household member is given by the variable *relate*. The list of possible values that this variable can take on changes over time. I classify an individual as being the head's spouse in the years 1962-1994 if, simply, the individual is classified as "spouse" in the survey. From the year 1995, I consider the individual as a spouse if its relationship to the household head is categorized as either "spouse" or "unmarried partner", as the latter was added from the year 1995. I exclude all same-sex couples. Before the year 1995, there are no same-sex couples in the data, while after that they are few in number.

I group households based on the levels of education of the two individuals in a couple have, or by sex and the level of education of the individual in a single household. Importantly, I, e.g., differentiate between a household in which the woman is low skilled and the man in high skilled, and a household in which the woman is high skilled and the man is low skilled. For practical reasons, I always classify the man in a couple to be the household head. The only time at which this has any implications is when I classify households to be working age households, which I do based on the age of the head. Throughout this paper I denote low skilled by L, high skilled by H, men by m, and women by f. In couples, the different types are LL, LH, HL, and HH. The first capital letter indicates the skill level of the man, and the latter indicates the skill level of the woman. The four types of single households are Lm, Lf, Hm, and Hf. E.g., Hf denotes high skilled women.

For practical reasons, an individual is classified either as low skilled or high skilled, although the data allows for finer groups. An individual is considered to be high skilled if he or she has at least a college degree, while the remainder are considered to be low skilled. As older cohorts, with relatively lower levels of education, are replaced with newer cohorts, with on average higher levels of education, this has led the fraction of high skilled individuals to increase since 1962, as is clearly visible in Figure 1.1. Over this period, the share has quadrupled, and the yearly increase has been close to 0.5 percentage points on average.

I focus on how individuals in the different types of households differ in the cross section and in terms of changes over time. How the decisions and outcomes of the individuals in the different types of households then map to aggregate quantities depends on the shares of households that are of each type. There are eight types of households: four types of single households, and four types of couple households. In each year, I compute the shares of households of each type. These are graphed for couples and singles in Figures 1.2a and 1.2b, respectively. By construction, the shares sum to one in each year.

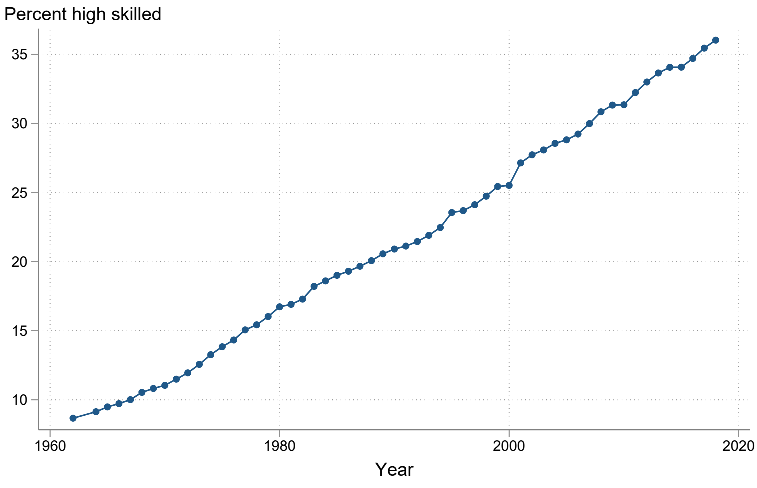

Figure 1.1: *Share of the adult population that is high skilled*

Note: The sample considered here consists of all individuals in the CPS that are aged 25 years or older for whom I observe information about education. See Section 1.2.1 for more details.

In addition to the increase in the average level of skills that was observed in Figure 1.1, the shares in Figure 1.2a are also affected by (i) the fraction of couples, (ii) the skill levels of individuals in couples, and (iii) sorting patterns. One of the more striking patterns, visible in Figure 1.2a, is that the fraction of couples in which both the man and the woman are low skilled falls dramatically over time. In part, this is driven by the fact that the fraction of households that are couple households has decreased over time.[1] The major driver behind

[1]The fraction of households that I categorize as couple households falls by around 20 percentage points over the period. Most of this decline happens until the beginning of the 1990s. See Figure 1.2a in Appendix 2.B. Greenwood et al. (2016) and Eckstein and Lifshitz (2011) show that this decline partly comes from decreased marriage rates, partly from increased divorce rates.

the remaining part of the decline is that there are fewer low skilled households among all households. Figure 1.2a also shows that the fraction of households that consist of two high skilled individuals has become considerable, whereas in the year 1962 they were rare.

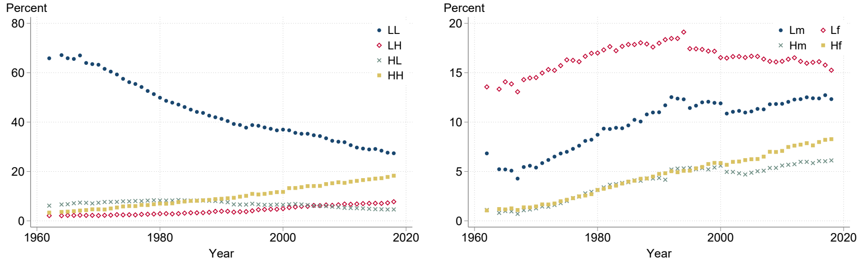

(a) *Couples households* ***(b)*** *Single households*

Figure 1.2: *Fractions, by type of household*

Note: *Panel(a)*: Shares of different types of couple households. The first letter in the legend indicates the skill level of the man and the latter the skill level of the woman. *Panel(b)*: Shares of different types of single households. The first (upper case) letter in the legend indicates the skill level of the individual, and the latter (lower case) letter indicates the sex.

Something that is not apparent from Figure 1.2a is the sorting patterns in couples. Figure 1.3a shows that the fractions of men with high skilled spouses have increased over time for both low skilled and high skilled men. However, while the increase in percentage point terms was greater for high skilled men, the increase in relative terms was greater for low skilled men, as is illustrated in Figure 1.3b, which plots a bias ratio. One interpretation of this ratio is that it captures how many more times likely it is, at some given year, to randomly select a high skilled woman out of the pool of women in couples with high skilled men, compared to how likely it is in the pool of women in couples in which the men are low skilled. If this result is interpreted as a decrease in assortative mating, it goes in the opposite direction of what is found in Greenwood et al. (2016). The difference in results is likely to depend on differences in methods, not differences in data. In Appendix 1.A.1 I discuss the differences in results and methods, and also present another interpretation of the bias ratio.

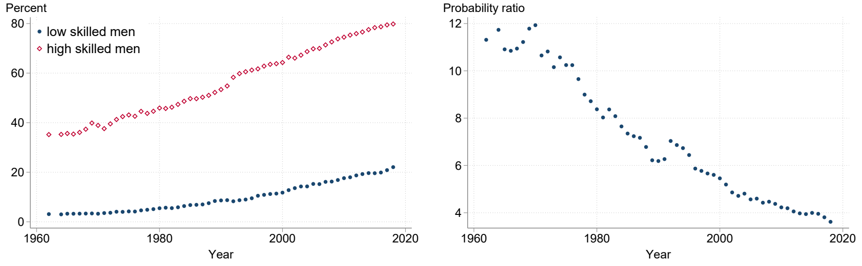

(a) *Men with high skilled spouses* ***(b)*** *Probability ratio*

Figure 1.3: *Sorting*

Note: *Panel (a)* shows the fractions of low skilled and high skilled men whose spouse is high skilled. Circles in *panel (b)* are created by, in each year, dividing the red diamonds by the blue circles from panel (a). The probability ratio can be interpreted as how many more times likely it is, at some given year, to randomly pick a high skilled woman out of the pool of couples with high skilled men, compared to how likely it is in the pool of couples in which the men are low skilled. I did also construct the corresponding figures but with cohort fixed effects, to be able to adjust for the fact that, at any given point in time, low skilled men tend to be older than high skilled men, and therefore low skilled men are more likely to have low skilled spouses, as the pool of low skilled spouses is higher. The graphs are essentially indistinguishable from the two panels in this figure and therefore I do not show them.

Over time, the fraction of high skilled individuals increased more in the group of couple households than among singles (Figure 1.4). In the year 1962, the fraction of high skilled individuals in single households exceeded the fraction of high skilled individuals in couple households, both among men (Figure 1.4a) and women (Figure 1.4b). But since around the year 2000, a higher fraction of men in couples, than single men, are high skilled, and the differences still seem to be diverging. The picture is similar for women, with one difference being that the fraction of high skilled women has been higher among women in couples since around the year 1990.[2]

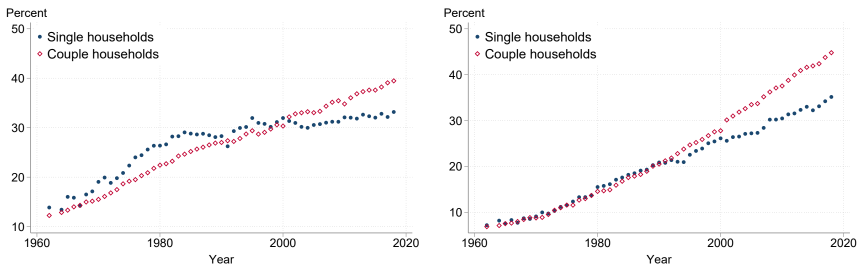

(a) *Men* ***(b)*** *Women*

Figure 1.4: *Fraction high skilled: singles vs individuals in couples*

Note: *Panel (a)* shows the share of high skilled men in single households (blue circles) and couple households (red diamonds). *Panel* (b) illustrates the same development, but for women. Adjusting for cohort-fixed effects has minor effects on the results, so I do not plot the graphs. Data comes from the Current Population Survey (CPS) and I restrict it to individuals in households in which the head is within the age range 25-64. See the main text for more details.

Let me summarize what I have shown. First, the share of high skilled individuals has grown steadily and considerably since the year 1962. The largest contributor behind this increase is couples in which

[2]Eckstein and Lifshitz (2011) decompose married women into five groups based on their levels of education. They show that it is especially among women with (i) a high school degree, (ii) some college, or (iii) a college degree that the labor force participation rates increased. The increases are smaller for women that are high school dropouts and for those that have a post college degree.

both spouses are high skilled. Moreover, the shares of high skilled single men and women, respectively, especially the latter, have increased considerably since the year 1962. Regarding how individuals, based on their skill type, are sorted across single and couple households, I showed that there has been an increased tendency for high skilled men and women to be observed in couple households. Moreover, among couple households, the degree of positive sorting on skills, captured by the bias ratio shown in Figure 1.3b, has become significantly smaller.

1.2.2 Hours of Market Work

I rely on information from two types of variables in the CPS to calculate market hours. First, I use answers to the question about how many hours the respondent worked last week, *ahrsworkt*. A limitation is that it does not contain any information about how many weeks per year the individual worked, which of course is important when analyzing labor supply (and calculating hourly wage rates). Starting in the year 1976, individuals were asked how many weeks per year they work, *wkswork1*. This question was not asked in the years before that. However, some individuals were asked about an interval for the number of weeks that they worked (*wkswork2*).[3] I impute values for weeks worked for individuals for whom answers exist about the intervals, and then assume that the year-times-household type level averages for weeks worked are the same for individuals without any information about weeks worked.[4]

[3]The intervals are: 1-13; 14-26; 27-39; 40-47; 48-49; 50-52.

[4]In the first step, I select individuals from the years 1976-1980 for whom I see information about both the actual weeks and the weeks interval. Then weeks are regressed on the weeks intervals, where each interval enters as an indicator. Weeks are then predicted for individuals for whom I see the weeks intervals in the

Data on hours worked for individuals in the armed forces are not available and the variable is originally coded as NIU (not in universe), except in the year 1962 when it is coded as 0. I code the hours variable as missing for individuals in the armed forces in the year 1962. Individuals who are unemployed are usually coded as NIU, but I choose to instead assign them 0 hours of market work. I do this also for individuals who are not in the labor force. My reason for making these adjustments in the hours variable for unemployed individuals and individuals who are not in the labor force is to include effects along the extensive margin. Ramey and Francis (2009) show that the trends in market hours are significantly different depending on whether hours worked are compared with the number of individuals in the labor force or the working age population.

Before I proceed and analyze how market hours differ in the cross-section and have developed over time, I exclude some observations based on income. I start by including income from three sources: (i) income that the household received as an employee (*incwage*), (ii) non-farm business income (*incbus*) and, (iii) farm income (*farminc*), either as a tenant farmer or income from operating a farm that the individual owns. Denote this sum by *labinc*. By dividing this sum by the number of hours worked I have a measure of the *gross* hourly wage rate. I trim the data by excluding observations where (i) *labinc* is negative, or (ii) the hourly wage rate is in the top percentile in the hourly wage distribution.[5]

Figure 1.5a shows that the average number of market hours in

years prior to 1976. Lastly, I calculate the average number of weeks worked for individuals in each household type in the years prior to 1976, and assume that weeks worked are the same for those individuals where the weeks interval does not exist.

[5]This is done year-by-year.

the households that I consider (i.e., classified as head or spouse in a household in which the head is in the working age bracket) increased by almost three hours between the years 1962 and 2018. The changes in market hours are different across many dimensions, and this section focuses on some of these. The first thing that can be established, by looking at Figure 1.5b, is that the increase is explained by the fact that women, as a group considered, increased the number of hours that they worked in the market, while the market hours for men decreased.

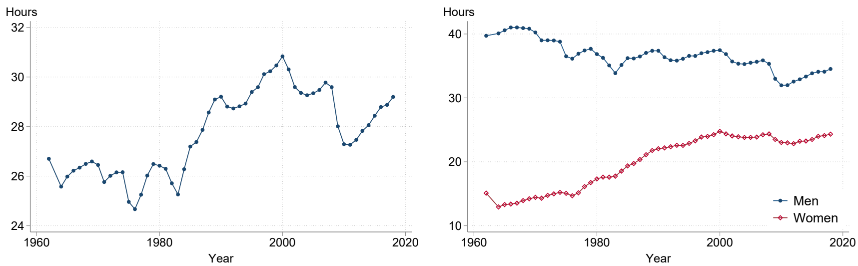

(a) *Average* ***(b)*** *By sex*

Figure 1.5: *Market hours*

Note: *Panel (a)*: average number of hours worked in the market. Men and women are pooled; *Panel (b)*: average number of hours worked in the market among men and women, respectively.

A comparison with Ramey and Francis (2009) indicates that these results for average hours worked are similar to theirs. I compare average hours for some selected years, and the results are shown in Table 1.1. In any given year, in any given group, the differences between average hours computed here and average hours computed by Ramey and Francis are small. Moreover, the changes over time, for each respective group, are similar. E.g., while I find that average hours worked increased by 4.4 between the years 1980 and 2000, they find the increase to have been 4.1. For men, we both find the increase over this period to have been around one hour, while for women it was 7-8 hours.

Table 1.1: *Comarison with Ramey and Francis (2009)*

	(1) Both sexes Here	(2) Both sexes RF	(3) Men Here	(4) Men RF	(5) Women Here	(6) Women RF
1960/1962	26.7	27.0	39.7	40.8	15.1	13.9
1980	26.4	28.0	36.8	36.3	17.3	20.0
2000	30.8	32.1	37.5	37.6	24.8	26.8
2005	29.3	31.3	35.5	36.8	23.8	26.1

Note: This table compares average hours worked that are computed in this paper with average hours worked from Ramey and Francis (2009) (see their Table 2), for some selected year. Note that the first comparison is done between two different years, the years 1960 and 1962, as Ramey and Francis (2009) use data from 1960, while I start in the year 1962. Moreover, note that the groups considered here are not the same as in Ramey and Francis (2009). Two important differences are that: (i) Ramey and Francis (2009) have information about hours worked in the military while I do not, and (ii) while I consider the age group 25-64, they pool individuals in the age group 25-54.

It is a well-established fact that labor force participation among women increased dramatically in the post-war era, especially among women in couples (see, e.g., Costa, 2000; Eckstein and Lifshitz, 2011; McGrattan et al., 2004). Various channels behind this increase have been proposed. E.g., Eckstein and Lifshitz (2011) find increased education of women to be the most important factor, while Greenwood et al. (2005) and Greenwood et al. (2016) argue that a decreasing price of capital (used in home production) was more important. Figure 1.6 illustrates the employment rates among women, measured as the share of women with positive market hours, in single and couple households, respectively. While the share grows by almost 10 percentage points for single women, it doubles among women in couples. McGrattan et al. (2004) also find this stark difference in developments for single women and women in couples, respectively. Eckstein and Lifshitz (2011) investigate the labor force participation among married women by cohorts. Cohorts born during the 1960s participated to a much greater extent in market work than did earlier cohorts. The increase ceased after that.

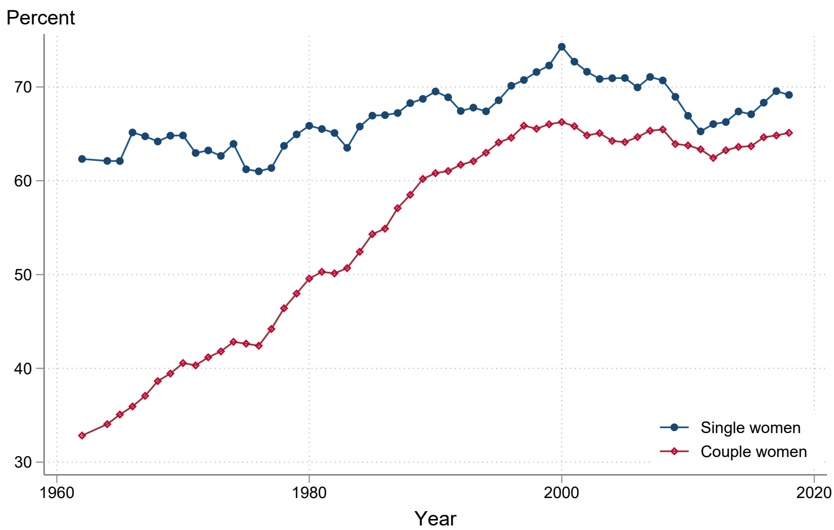

Figure 1.6: *Employment rates among women, by status*

Note: This figure shows the fraction of women in single households (blue circles) and couple households (red diamonds), respectively, with positive hours of market work. I did also create the corresponding graph but with age fixed effects. Although the average age of women in the two respective groups develops somewhat differently over the time period, the graph that includes age fixed effects is very similar, indicating that the different paths in average age cannot account for the differences in this figure.

Figure 1.7 shows that hours worked in the market by women differed not only depending on whether the woman had a partner or not, but also by skill level. Figure 1.7a shows that high skilled single women worked considerable more hours than low skilled single women, and that this difference was greater in the year 2018 than it was in the year 1962.[6] As is seen from Figure 1.7b, there were considerable differences in how many hours women in couples worked, depending on their own level of education as well as on their partner's level of education. E.g., while high skilled women with low skilled partners, on average, worked more than 20 hours in the market in the year 1962, high skilled women with high skilled partners worked only about 15 hours in the market. To a large degree, the cross-sectional differences that can be observed for the year 1962 persisted up until the year 2018. Still, the time-series developments have been different. It was especially low skilled women with high skilled partners who increased their labor supply, but clearly there have been significant increases in all four groups.[7] [8]

The employment rate among men in couples was higher than among single men in every single year since the year 1962 (see Figure 1.8). In contrast to what was observed for women, disregarding shorter-run movements, there were no differences in terms of changes over time: the employment rate dropped by around five percentage points in both groups.

Dividing men into finer groups exposes systematic cross-sectional

[6]That the difference is greater in 2018 can be seen more clearly in Figure B.2a in Appendix 2.B.

[7]While some of the increase in hours is due to increases along the intensive margin, a large majority can in each group be attributed to increased employment rates.

[8]Go to Figure B.2b in Appendix 2.B to see differences in time trends more clearly.

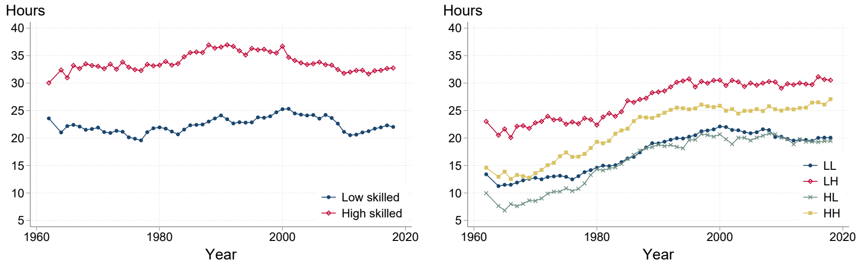

(a) Single women *(b) Women in couples*

Figure 1.7: *Market hours among women in different types of households*

Note: *Panel (a)* Average hours worked among two groups of single women: low skilled (blue circles) and high skilled (red diamonds). *Panel (b)*: Average hours worked among women in four different kinds of couple households. The legend in the graph tells the combination of the woman's as well as the man's skill level. The first letter indicates the skill level of the man and the latter the skill level of the woman; e.g., LH (red diamonds) represents the development of average hours worked in the market for a woman who is high skilled and whose partner is low skilled.

differences in average hours worked in the market, but very similar developments for these hours over time. Similar to how it was for women, high skilled single men have been working about 10 hours more per week in the market than low skilled single men, as can be seen in Figure 1.9a. This difference has been stable over time: average hours declined after the year 1962 by similar magnitudes in both groups (see this more clearly in Figure B.3a in the appendix).

The number of hours that men in couples worked depends on the combination of their own skill level and their partner's skill level (see Figure 1.9b). High skilled men with high skilled partners worked the most hours. At the bottom we find that low skilled men with low skilled partners worked the least amount of hours, on average. In

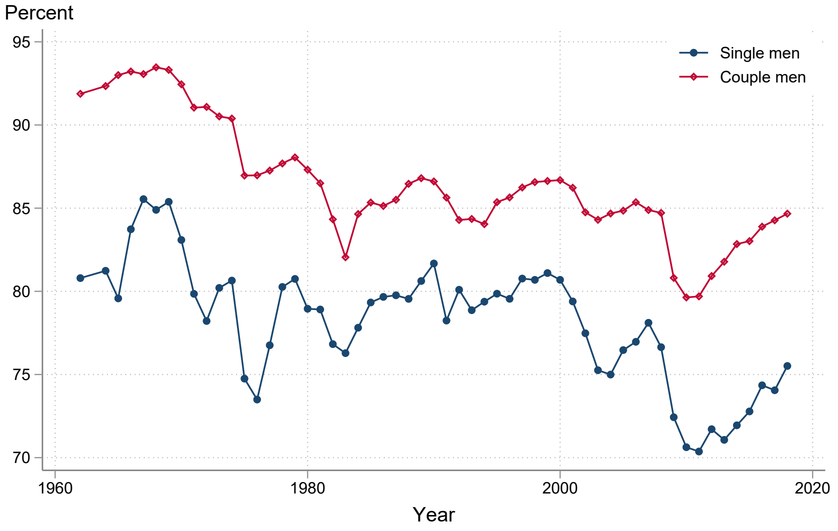

Figure 1.8: *Employment rate among men, by status*

Note: This figure shows the fraction of men in single households (blue circles) and couple households (red diamonds), respectively, with positive hours of market work. I did also create the corresponding graph but with age fixed effects. Although the average age of men in the two respective groups develops somewhat differently over the time period, the graph that includes age fixed effects is very similar. I did also look at the extensive and intensive margin separately. In terms of changes over time, in neither of these cases can any notable differences between men in single and couple households be observed.

each of the four types of couple households, the changes over time are very similar. The only group for which average hours worked in the market potentially dropped somewhat more is low skilled men with low skilled partners. But it can be seen in appendix Figure B.3b that the differences are relatively small.

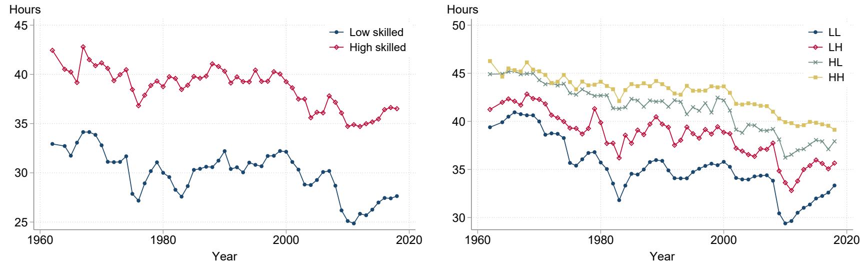

(a) *Single men* ***(b)*** *Men in couples*

Figure 1.9: *Market hours among men in different types of households*

Note: *Panel (a)* Average hours worked among two groups of single men: low skilled (blue circles) and high skilled (red diamonds). *Panel (b)*: Average hours worked among men in four different kinds of couple households. The legend in the graph tells the combination of the man's as well as the woman's skill level. The first letter indicates the skill level of the man and the latter the skill level of the woman; e.g., LH (red diamonds) represents the development of average hours worked in the market for a man who is low skilled and whose partner is high skilled.

1.2.3 Expenditures

For expenditures data at the household level, I use the Consumer Expenditure Survey (CE Survey, or CEX). I start with the so-called *MTBI files* from the year 1980 up until and including the year 2018. The files contain detailed information about monthly expenditures at the household level. The MTBI files are included in the broader set of *interview files*, which are all a result of the *Interview Surveys*. In the monthly expenditures files, expenditures are categorized by a Universal Classification Code (UCC). For my purposes, I need expenditures categorized into the three major categories: services, nondurable goods and durable goods, respectively. For most UCCs, I map them into these categories by using the crosswalk from Coibion et al. (2021). I complement the crosswalk with new UCCs that have been added since the year 2017. See Appendix 1.A.8 for more details. For each household, I calculate the share of expenditures that is allocated to each major category.

After having grouped expenditures into the three spending categories, and households into one of the eight respective household types (following the same scheme as in Section 1.2.2), I summarize them by the mean expenditure shares, in each major expense category, for each combination of year and household type.

These are not the final expenditure shares, however, as I make significant adjustments to them. There are two reasons for these adjustments. First, I adjust the expenditures in order to line up with aggregate National Income and Product Accounts (NIPA) statistics. Second, I adjust expenditures on durable goods/capital to better reflect the user cost of capital, rather than reflecting the expenditures on new durable purchases. I go through these procedures in detail in Appendix 1.A.8.

It is a well-established fact that, over time, households have allocated an increasing share of total expenditures to service consumption. Moreover, it has been documented, e.g., by Boppart (2014), that this expenditure share is higher among households with larger incomes. Thus it is not surprising that the expenditure shares on services differ between household types, see figures 1.10a and 1.10b. What might come as somewhat of a surprise, though, is that, holding the skill levels fixed, the expenditure shares on services are very similar when comparing single men to single women.

Figures B.4 and B.5 in Appendix 2.B show how expenditure shares on nondurable goods and durable goods, respectively, developed in the different types of households. The decrease in the expenditure shares on nondurable goods was substantial for all household types. Single

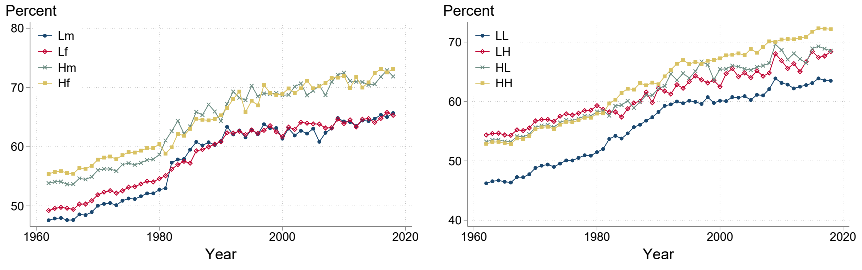

(a) Single households *(b) Couple households*

Figure 1.10: *Expenditure share on services*

Note: *Panel (a)*: Expenditure shares on services among single men and women, depending on skill levels. *Panel (b)*: Expenditure shares on services among couple households, depending on levels of education of the two spouses.

women spent a significantly larger fraction of their budget on nondurable goods than single men. Among couples, households in which both spouses were low skilled spent a larger fraction on nondurable goods than did other couple households. The differences were small between the other couple households. Between the years 1962 and 2018, the expenditure shares on durable goods fell. The decrease was substantial for single men, but modest for single women, and couple households. In the later years, the expenditure shares on durable goods were similar for all household types.

Equation (1.1) shows how to calculate an index for the quantity ratio between expenditure types i and j using expenditure shares and prices:[9]:

$$q_{ij} = \frac{\eta_i}{\eta_j}\frac{P_j}{P_i} \tag{1.1}$$

In Figures 1.11a and 1.11b, I show how this quantity ratio between services and nondurable goods developed for different types of single and couple households. Not surprisingly, the ordering of household types in each respective panel closely resembles the ordering when the focus was on service expenditure shares, in Figure 1.10. In terms of developments over time, a dominant trend is that the ratio increases up until the 1990s. Thereafter, developments were somewhat heterogeneous across household types. For most household types, the ratio then developed sideways, or fell somewhat. In couple households in which both partners are high skilled, the ratio increased further after the 1990s.

1.2.4 Wage Rates

I now proceed with an investigation into hourly wage rates. Note, first, that earnings are deflated by the service price index. Dividing the income variable *labinc* (defined in Section 1.2.2) by hours worked yields *gross* hourly wage rates. I am, however, interested in an hourly rate that relates more closely to how the labor choice affects disposable income, which means taking taxes and transfers into account. Therefore, *labinc* (denoted by y_{pre}) is mapped to disposable income (denoted by y_{post}), by using the tax-and-transfers function from Heathcote et al. (2017). For details about this mapping, see Appendix 1.A.3.

[9]Data on prices are from the NIPA. All prices are normalized with respect to the service price index. I briefly introduce the data and plot time series for the relative prices of nondurable and durable goods, respectively, in Appendix 1.A.2.

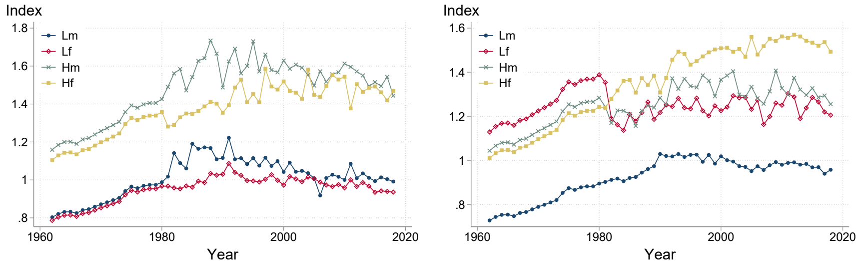

(a) *Women* ***(b)*** *Men*

Figure 1.11: *Ratio of services vs nondurable goods*

Note: *Panel (a)*: Ratio of services vs nondurable goods for different types of single households. *Panel (b)*: Ratio of services vs nondurable goods for different types of couple households. Series in both panels are indexed to the value of the ratio for low skilled single men in the year 2010.

By dividing disposable income by hours worked I get the (disposable) hourly wage rate.

Two issues are worth mentioning related to households' disposable incomes. First, interest income and dividend income are abstracted from. For all household types, the income from returns on capital is small in relation to labor income.[10] Therefore, an introduction of capital income would probably have small effects. My other reason for excluding it, is that I want to map the model in Almgren (2023) to the data. How to alter the model in order to incorporate unearned income

[10]Using CPS data from the years 1968-2018, I can investigate how large capital income (call it *capinc*) is in relation to labor income for men and women in different types of households. I include the following (IPUMS) variables as capital income: years 1968-1975: *incidr*; years 1976-1987: *incdrt* and *incint*; years 1988-2018: *incint*, *incdivid* and *incrent*. Then I summarize *capinc* and *labinc* by their mean values in each household type-year combination. Finally, I take the ratio of these two. For more than 92 percent of these type-year combinations, the amount of capital income in relation to labor income is less than 10 percent. The unweighted average is 5.5 percent.

is not obvious, and would lead to less tractable results. Second, there are two obvious issues with the tax-and-transfers function and how it is used. Concerning the function itself, it imperfectly maps gross to disposable income (Heathcote et al., 2017). The other issue is that I do not explicitly use the function in the model, but instead incorporate it in an indirect way when I calculate (disposable) wage rates, which are then taken as given by the households. This ignores the connection between gross wage rates, the labor choice, and marginal tax rates, but greatly simplifies the solution.

Figure 1.12a shows that, between the years 1962 and 2018, nominal wage rates grew more than the price of services (i.e., real wages, as I sometimes will refer to them as henceforth, increased), for both men and women. The figure also clearly shows a significant wage gap between men and women. This gap decreased over time, however, as the wage growth of women has considerably outpaced the wage growth of men since the early 1980s. While real wages did increase by around 40 percent over the period for men, real wages for women increased by an additional 20 percentage points (see Figure 1.12b). These results align with what is found in Eckstein and Lifshitz (2011), who find that the ratio between the mean wage rates of men and women was roughly constant into the 1980s, but thereafter decreased considerably.

Wages grew very similarly over the long run for women in couple households and women in single households (see Figure 1.13). The increases since the year 1962 amount to around 60 percent.

Cutting the women into finer groups exposes significant heterogeneity. From Figure 1.14a, it is easily seen that the skill premium was significant over the whole period. Comparing the years 1962 and 2018, the skill premium was lower in the latter of the two years, and, as can be more easily seen in Appendix Figure B.6a, it varied signifi-

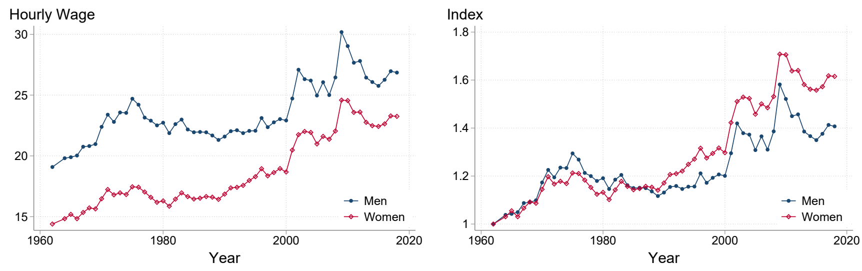

(a) *Mean wage for men and women* ***(b)*** *Indexed series*

Figure 1.12: *Hourly Wage Rates*

Note: *Panel (a)*: Shows mean wages, by year, for men and women, respectively. Nominal wages are deflated using the price index for services, for which the base year is set to 2010. *Panel (b)*: Indexed versions of the series in panel (a).

cantly over time. It is clearly visible that the skill premium for women dropped substantially in the 1970s, to then bounce back.[11]

There were also differences in wage levels and their growth rates for women in couples, depending on their own and their partner's level of education (see figures 1.14b and B.6b). Wages among high skilled women with high skilled partners (HH) were typically the largest, while low skilled women with low skilled partners had the lowest wages. Since the year 1962, the wage rates increased considerably among all four types of women. While the real wage rate increased by around 40 percent for high skilled women with low skilled partners, the increases were around 30 percent for the rest.

Average wages among men in couple households have been greater than among men in single households (see Figure 1.15a). Over time,

[11]In Appendix 1.A.4 I show that different trends in the age composition explain only a small part of the differences in hourly wages between high skilled and low skilled single women.

Figure 1.13: *Hourly wage rates, by status*

Note: Shows average wages, by year, single women and women in couple households, respectively. Nominal wages are deflated using the price index for services and the base year is set to 2010.

the wage rates among men in couples have increased more than among single men, as is clear from Figure 1.15b. However, in Appendix 1.A.5 I show that this is driven by heterogeneous developments in the age composition. Had developments in the age composition been the same in the two groups, the relative difference in wage rates would instead have been around 10 percent both in the year 1962 and the year 2018.

That there has been a substantial skill premium also among men is apparent from Figures 1.16a and 1.16b. What also becomes very clear by looking at the respective panels in Figure B.2a is that the skill premium is higher in the year 2018 than it was in the year 1962.

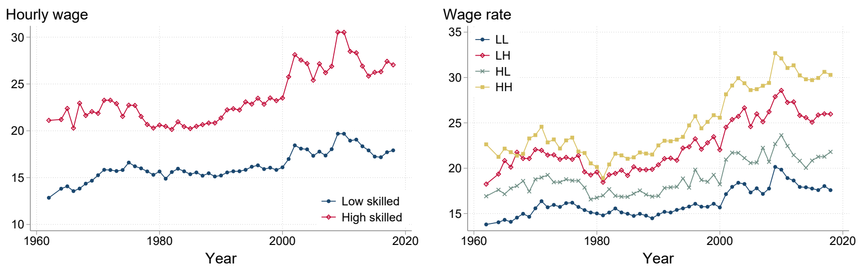

(a) *Single women* ***(b)*** *Women in couples*

Figure 1.14: *Wage rates for women in different types of households*

Note: *Panel (a)*: Median wage rates among two groups of single women: low skilled (blue circles) and high skilled (red diamonds). *Panel (b)*: Median wage rates among women in four different kinds of couple households. The legend in the graph tells the combination of the man's as well as the woman's skill level. The first letter indicates the skill level of the man and the latter the skill level of the woman; e.g., LH (red diamonds) represents the development of average hours worked in the market for a woman who is high skilled and whose partner is low skilled.

Up until the 1990s, wages changed by about the same magnitudes for all types of men, irrespective of whether they were high or low skilled, or whether they had a partner or not (see Figures B.7a and B.7b Appendix 2.B). But from the 1990s, the skill premium increased dramatically, with the exception that high skilled men with low skilled spouses saw increases in their wage rates that did not exceed the increases among low skilled men by as much. Looking more closely at the data, the divergence occurred during the 1990s and the beginning of the 2000s, and has then stabilized.[12]

[12] An analysis in Appendix 1.A.6 shows that among single men, the increase in the relative wage among high skilled, compared to low skilled, is explained by the age composition.

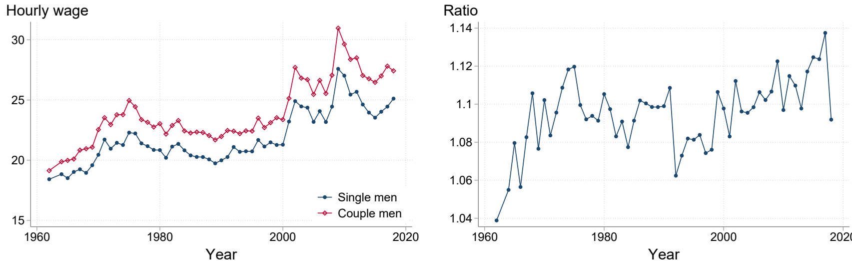

(a) *Wage rates, by status* ***(b)*** *Ratio*

Figure 1.15: *Hourly wage rates among men, by status*

Note: *Panel (a)*: Shows average wages, by year, for men in single households and couple households, respectively. Nominal wages are deflated using the price index for services and the base year is set to 2010. *Panel (b)*: Average wages of men in couple households relative to average wages of men in single households.

1.2.5 Hours Worked in Home Production

The current section documents hours worked in home production (I sometimes refer to them as *home hours*) in different types of households, and how they have developed over time. This section also serves as a motivation for why hours in home production should be considered. First, they are significant: both men and women, in all household types, spent a significant share of their non-leisure time in home production. Second, differences between men and women, and across different types of households, can be be substantial, and there have been significant developments over time.

For hours in home production, I use data supplied by Aguiar and Hurst (2007), which ends in the year 2003, and complement it with data up until the year 2018. Since the year 2003, the source of the time-use data is the American Time Use Survey (ATUS) and it is available at a yearly frequency. In the earlier years, data is available

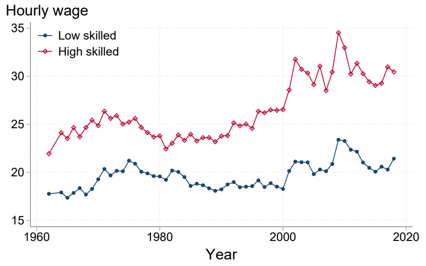

(a) Single men

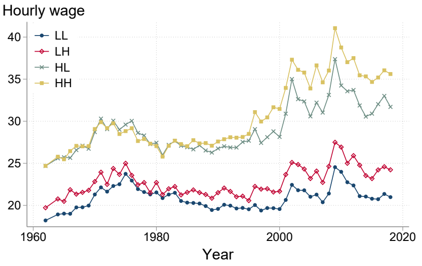

(b) Men in couples

Figure 1.16: *Wage rates for men in different types of households*

Note: *Panel (a)*: Average wage rates among two groups of single men: low skilled (blue circles) and high skilled (red diamonds). *Panel (b)*: Average wage rates among men in four different kinds of couple households. The legend in the graph tells the combination of the man's as well as the woman's skill level. The first letter indicates the skill level of the man and the latter the skill level of the woman; e.g., LH (red diamonds) represents the development of average hours worked in the market for a man who is low skilled and whose partner is high skilled.

at irregular intervals and are of different origins. More specifically, Aguiar and Hurst (2007) link data from four additional time-use surveys: 1965–1966 America's Use of Time; 1975–1976 Time Use in Economics and Social Accounts; 1985 Americans' Use of Time; and the 1992–1994 National Human Activity Pattern Survey.[13]

I follow Aguiar and Hurst (2007) when assigning activities to the category *home production* (which they refer to as *total nonmarket work*, see Table IX in their paper).[14] Households are either categorized as low or high skilled, in the same way that households were

[13]See Aguiar et al. (2012) for some more details about each survey.

[14]In Appendix 1.A.6 I compare the results for average hours worked in home production with the results in Ramey (2009), where, mainly to differences in how activities are categorized, some differences emerge.

grouped in the CPS. In the earlier years, the number of observations per household type is small and there is often one observation that heavily influences the mean. To limit the influence of these, I drop the observation with the largest value for hours in home production in each year, for each respective household type. As I have information about the other variables and prices at a yearly frequency (or have imputed), I will fill in the gaps in the years where I do not have information about hours worked in home production, by linear interpolations. I weigh the data using the sampling weights within each of the time-use surveys.

As with the other variables, I focus on individuals that are at least 25 years old but not older than 64. There is one severe limitation in the data: while I can see the educational level of the spouse from the year 2003 and onward, this information is not available in earlier years. I will look at hours in home production in the different types of couple households for the years in which the necessary data is available, but need to restrict the analysis to less detailed cuts in the earlier years. This data limitation is also a problem with regards to taking a model with couple households to the data.

Figures 1.17a and 1.17b illustrate the cross-sectional differences and time developments of hours in home production for men and women, respectively, in the different types of couple households. Clearly, for both men and women, there is a substantial year-to-year variation in the averages. I do not show it in the paper, but a short investigation into these changes reveals that the differences from one year to the next in a large majority of the cases are not statistically significant. That said, there are some differences and changes in the averages that are worth noting. Among women, there was a rather clear tendency for low skilled women to work more hours in home production that did high skilled women. A regression with women, in which I pool all the years, shows a statistically significant difference between the number of hours that both types of high skilled women worked in home production, compared to the women who are low skilled and have low skilled partners. The difference in means between the two types of low skilled women is, however, not statistically significant. The same type of investigation into the cross-sectional differences between different types of men's hours does result in some statistically significant results. Low skilled men with high skilled partners worked more in home production than did low skilled men with low skilled partners, while high skilled men with high skilled partners worked fewer hours in home production. In both cases, the differences are smaller than one hour.

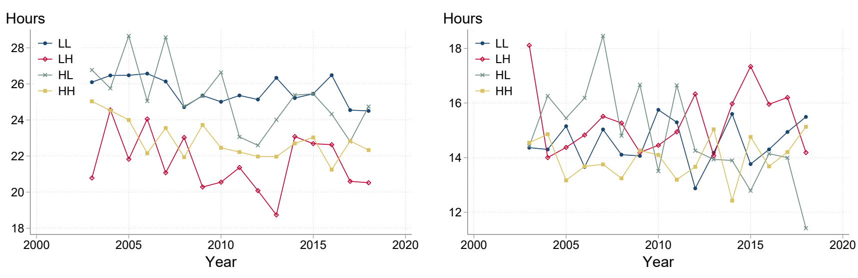

(a) Women *(b) Men*

Figure 1.17: *Hours worked in home production: couples*

Note: *Panel (a)*: Hours worked in home production for women in different types of couple households. *Panel (b)*: Hours worked in home production for men in different types of couple households.

Focusing on time-trends, Figure 1.17a suggests that women in all types of couple households decreased their hours in home production over the period that is considered here. To test the changes statistically, I group the data from the years 2003-2005 into period 1 and data from the years 2016-2018 into period 2. I then test if the hours worked in home production are statistically different in the second period compared to the first, across the different types of households. Average hours declined for high skilled women with low skilled partners, but this decline is statistically insignificant. For the three remaining groups, the decreases are statistically significant. Among the different types of men, the change is statistically significant only for high skilled men with low skilled partners.

When only considering each individual's own educational level, it is possible to create longer time series. I group individuals by gender and skill level, and by whether they live in a couple or single household, and then compute averages for each group in each year when data is available. As mentioned earlier, I trim the data by removing the largest value for each group-year combination. This only has notable effects on the averages in the earlier years, when the number of group-year observations is sometimes low.

The resulting time series are illustrated in the two panels of Figure 1.18. Women in couples spent more hours in home production than single women. Among women in couples, low skilled women systematically worked more in home production than high skilled women. It was also the case that low skilled single women worked more hours in home production, compared with high skilled single women. Concerning changes over time, we see noticeable decreases for three of the groups, with high skilled single women being the exception. The changes for women in couples are remarkable and amount to around -15 hours. Compared to in the year 1962, the differences between single women and women in couples in the year 2018 were significantly smaller.

Between the groups of men, the cross-sectional differences were relatively small in any of the years. A tendency for single men to have worked more hours in home production in the earlier years can be observed. Over time, there have been noticeable, albeit modest, increases. Hours worked in home production increased more among men in couple households than among single men, which led to men in couples working relatively more in home production in most years after the year 1980.

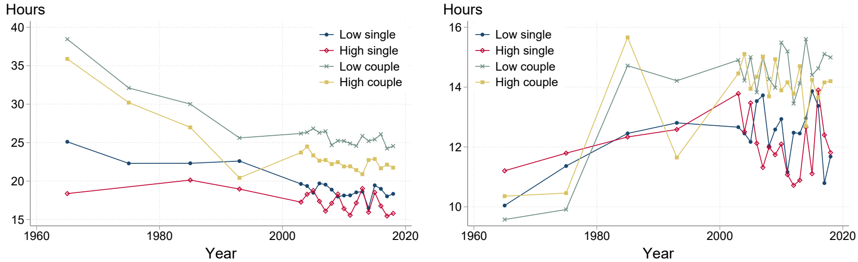

(a) *Women* ***(b)*** *Men*

Figure 1.18: *Hours Worked in Home Production*

Note: *Panel (a)*: Hours worked in home production for women, depending on their skill levels and whether they live in a couple or single household. *Panel (b)*: Hours worked in home production for men, depending on their skill levels and whether they live in a couple or single household.

Changes over time in the age-composition in the different groups seem to have contributed to how the plain averages have developed. In Appendix 1.A.7, I show that the age-adjusted means, for men, have risen more, but that qualitatively the results are very similar. For women, the differences are negligible.

It has been shown that both women and men, in all types of households, spent a significant portion of their time in home production.

I have also presented time series that show that especially women in couples decreased the numbers of hours they work in the household. Men, on the other hand, have increased the number of hours that they work in home production. Comparing the two panels in Figure 1.18, it is, however, quite clear that the total hours in home production that are worked by men and women in couples has fallen. It seems plausible that individuals' choices about how much to work in the market, and how much to work in the household, respectively, are not two independent problems. Rather, these choices are made simultaneously. Many things may affect how they interact, however; e.g., (i) what is produced at home and how, and (ii) to what extent can home produced output substitute market produced output? Developing an explicit framework, with home production, to address these questions is the focus in Almgren (2023).

1.3 Hours Worked in the Market and Wage Rates: A Comparison Between Women and Men

Some statistics that can be inferred from the already presented data series, but deserve to be analyzed explicitly, are the relative quantities of men's and women's hours, and wage rates, respectively, in different types of couple households. I start by considering market hours.

Figure 1.19a shows the number of average hours that women worked in the market compared to men, in each type of couple household that is considered. Two things become very clear: (i) there is substantial heterogeneity, and (ii) all the ratios have grown considerably over the period. In terms of heterogeneity in levels, we see that the ratio is higher by a large margin among couples in which the woman is high skilled and the man is low skilled. In the year 1962, women in these types of households worked almost 60 percent of the hours that their men worked in the market. By the year 2018, this percentage exceeded 80 percent. In Figure 1.19b, I illustrate the relative changes in each of these ratios. The relative changes in these ratios are inversely related to the levels: market hours among women in households in which women worked relative little compared to men have seen larger positive changes since the year 1962. The relative number of market hours that low skilled women worked compared to their partners who are high skilled did more than double.

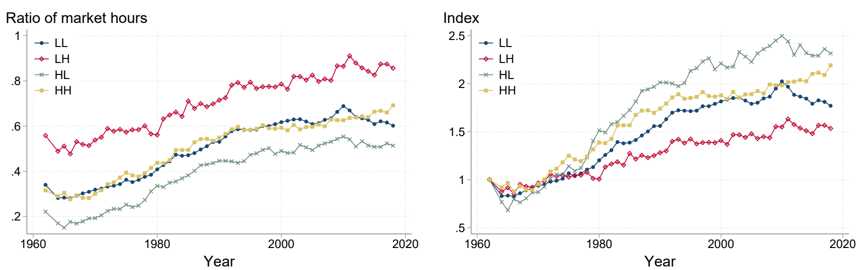

(a) *Ratio* ***(b)*** *Index*

Figure 1.19: *Women's market hours relative to those of men*

Note: *Panel (a)*: Shows relative market hours, by year, for women compared to men, in different types of couple households. *Panel (b)*: Indices of the data series plotted in panel (a).

Another statistic of interest is women's wage rates relative to those of men in different types of couple households. As naturally was to be expected, the ratio was highest in couples where the woman is high skilled and the man is low skilled, and lowest in the opposite case (see Figure 1.20a). In the later years, the ratios were similar in couples in which both spouses had the same skill level. The relative wage of women in households where both were low skilled was, however, at lower levels during the first four decades. The developments of the relative magnitudes are illustrated in Figure 1.20b. Comparing the years 1962 and 2018, the changes were overall modest. While the relative wage rates of women with high skilled partners decreased little or stayed close to the initial levels, the relative wages of women with low skilled partners increased by around 10-15 percent.

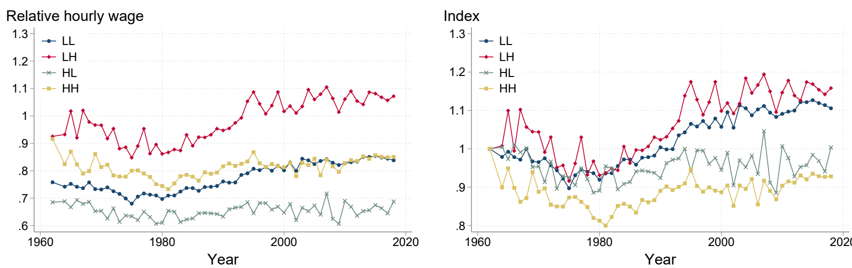

(a) *Ratio* ***(b)*** *Index*

Figure 1.20: *Women's wage rates relative to those of men*

Note: *Panel (a)*: Shows relative hourly wage rates, by year, for women compared to men, in different types of couple households. *Panel (b)*: Indices of the data series plotted in panel (a).

In particular two observations are worth mentioning from a comparison between women's and men's relative wage rates with women's and men's relative market hours. First, the ordering of the relative wages is the same as the ordering of the relative hours across the different types of couples, in all years (disregarding that LL and HH households switch places from time to time). Second, it seems like relative market hours grew less in households where relative wage rates grew more (households with low skilled men).

Another possible comparison is between single men and women. Making the comparison within skill-types shows that single women had lower wage rates than men. The difference between men and women has historically been greater in the low-skilled group, as is clearly visible in Figure 1.21a. Between the years 1962 and 2018, the ratio of wage rates between high skilled women and high skilled men has fluctuated around 0.9 but there is no visible trend. The wage rates of low skilled women, however, increased relative to the wage rates of low skilled men; while the wage rates of low skilled women were on average 78 percent of the wage rates of low skilled men in the 1960s and 70s, the share averaged 86 percent after the year 2000. Figure 1.21b shows that market hours worked by both types of single women increased compared to those of their male counterparts. Jointly analyzing the two panels of Figure 1.21, it is particularly interesting to see that high skilled women started working more relative to high skilled men but that the relative wage rate between the two groups did not increase. This raises a question about what else, if not a change in the wage gap, that might be driving the relative hours gap upwards over time for high skilled women.

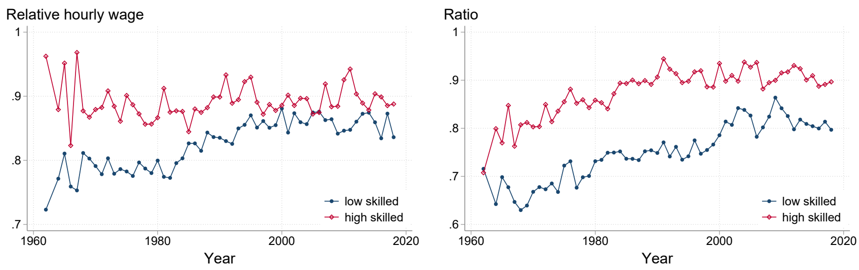

(a) *Relative wage rates* ***(b)*** *Relative market hours*

Figure 1.21: *Women's wage rates and hours relative to those of men*

Note: *Panel (a)*: Shows relative wages, by year, between women and men in different types of single households. *Panel (b)*: Shows relative hours, by year, between women and men in different types of single households.

1.4 Evaluating Theory

I will describe and evaluate the performance of a standard economic model, with stable homothetic preferences, in terms of its ability to match data on market hours and expenditures. The evaluation is restricted to households with single men.

The model is static and the production side of the economy is not modeled. Hence all prices are taken as given. The household chooses how many hours to spend working in the market, L, and the remaining hours are devoted to leisure. Working in the market generates disutility. The market wage rate is W and is type-specific.

Given its labor choice and resulting income, the household makes the decision about how to distribute expenditures. The simplest framework would not consider how expenditures are distributed but rather bunch everything into a consumption index C. This paper explicitly focuses on how expenditures are allocated between three main categories: *services*, *nondurable goods*, and *capital*, respectively. By necessity, to make the mapping possible between theory and data, the theory must separate between these types of consumption. Households have preferences over the three types of consumption, specified by a stable homothetic CES function. The problem is represented by the following constrained maximization problem:

$$\max_{M,N,K,L} U = \frac{1}{1-\gamma}\left(\omega_m^{\frac{1}{\sigma}} M^{\frac{\sigma-1}{\sigma}} + \omega_n^{\frac{1}{\sigma}} N^{\frac{\sigma-1}{\sigma}} + \omega_k^{\frac{1}{\sigma}} K^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma(1-\gamma)}{\sigma-1}} - \psi \frac{L^{1+\varphi}}{1+\varphi} \tag{1.2a}$$

subject to

$$WL = P_n N + P_k K + M \tag{1.2b}$$

where M, N, and K denote market services, nondurable goods, and durable goods (or capital), respectively. The wage rates and all prices are normalized with respect to the service price. Household types are not indexed here. Note, however, that wage rates are type specific and that this will lead to different outcomes for the choice variables for the different types of households.

1.4.1 Labor supply

According to the theory, hours worked by type s in year t are given by

$$L_{st} = \left(\frac{W_{st}}{P_t}\right)^{\frac{1-\gamma}{\gamma+\varphi}} \psi^{-\frac{1}{\gamma+\varphi}} \tag{1.3}$$

where $P \equiv \left(\omega_m + \omega_n P_n^{1-\sigma} + \omega_k P_k^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ is the implicit price index for the consumption index C.[15] As the price index is the same for both types of households, the theory's prediction about the between-household difference is a function only of the respective wage rates:

$$\frac{L_{high,t}}{L_{low,t}} = \left(\frac{W_{high,t}}{W_{low,t}}\right)^{\frac{1-\gamma}{\gamma+\varphi}} \tag{1.4}$$

From sections 1.2.2 and 1.2.4, respectively, we know that (i) high skilled single men worked more hours, and (ii) earned higher wages. A necessary and sufficient condition for this is that $\gamma < 1$ (not considering negative values of γ or φ), i.e., that the substitution effect dominates.

The market hours for both low and high skilled single men were considerably lower in the year 2018 than they were in the year 1962 (this was shown in Figure 1.9). The model suggests that hours in the year 2018 relative to in the year 1962, for type s, should be

$$\frac{L_{s,2018}}{L_{s,1962}} = \left(\frac{W_{s,2018}/P_{2018}}{W_{s,1962}/P_{1962}}\right)^{\frac{1-\gamma}{\gamma+\varphi}} \tag{1.5}$$

The time series for wage rates (see Section 1.2.4) and prices (see Appendix 1.A.2) show that the real wage increased between these two years, while the price index decreased. More to the point: W/P has increased for both types of single men households. To generate a decrease in hours worked, a necessary condition is $\gamma > 1$ and for the income effect to dominate. Clearly this simple theory cannot jointly explain cross-sectional and time-series evidence about hours worked.

[15] As prices and the wage rates are normalized with respect to the service index, the service index does not enter explicitly.

1.4.2 Allocation of Expenditures

Can the theory account for these facts? The quick and simple answer is: no it cannot. Regarding the cross-sectional differences within years, it is well known that a standard economic model with stable homothetic preferences generates homogeneous expenditure shares, and homogeneous quantity ratios. This is easy to see from an equation that shows the optimal quantity of services relative to nondurable goods:

$$\frac{M_{st}}{N_{st}} = \frac{\omega_m}{\omega_n} P_{nt}^{\sigma} \tag{1.6}$$

so if restricting the share parameters to be the same across households, there is nothing that can generate cross-sectional differences.

The relative price between nondurable goods and services has been falling steadily since the year 1962. From Equation (1.6) it is easy to see that the model prediction is then for the services-to-nondurable goods ratio to fall. This goes in the opposite direction of what is seen in the data betwee the year 1962 and the 1990s (see Figure 1.11). The model hence fails to match even the qualitative evidence on how households allocate their expenditures, both in the cross-section and over time.

1.5 Some Additional Observations

In addition to what was already discussed in Section 1.3, I want to highlight some additional observations from the data. Even though I do not here propose an explicit and well-defined theory, it is nevertheless easy to conclude that these observations are difficult to square with most standard extensions that one could make to the above proposed theory.

First, hourly wage rates have increased by roughly as much for single women as for women in couple households. This cannot be explained by differential trends in wage rates, as these have been quite similar. Here I do not propose anything specific, but conclude that the theory needs to treat single women separately from women in couple households.

Second, comparing high skilled single men with high skilled single women, it is easy to see that some patterns in the data call for model extensions. The market hours of high skilled single women have increased significantly over time but the relative hourly wage rate has not increased (this was shown in Figure 1.21). A similar observation is made by Eckstein and Lifshitz (2011) who observe that trends in schooling have been similar for men and women, but that the employment rates have gone in opposite directions.

In Section 1.4 I showed that the simplest framework cannot generate heterogeneity in expenditure shares. One way of generating that high skilled single men allocate a larger fraction of their expenditures to services than do low skilled single men, is with nonhomothetic preferences. If the wage elasticity of demand is higher for services than for goods, then this outcome can be generated. A framework that generates cross-sectional differences only via differences in wage rates would, however, falsely predict that the shares differ between single men and women, holding skill levels constant, since we know that their wage rates differ but not their expenditure share on services. Hence, something else is needed.

1.6 Conclusion

This paper documents cross-sectional differences and time-developments concerning time-use and expenditures. It has been shown that trends in hours worked in the market are different for men and women in different types of households. Moreover, hours worked in home production are significant, and they also vary across households and change substantially over time. The simple fact that hours worked in home production eat up a significant chunk of most individuals' time budgets should serve as enough motivation for considering it in economic models.

With respect to expenditures, the focus is on how they are allocated across market services, nondurable goods, and durable goods, respectively. I confirm one of the findings that is emphasized by Boppart (2014), that quantity ratios and relative prices move in opposite directions over many decades, which a standard model with stable homothetic preferences cannot generate.[16]

A contribution of my paper is that I separate individuals into groups based on a combination of gender, skill level, and whether the individual lives in a single or couple household. This exposes several interesting patterns. One of those is that the expenditure shares on services were very similar among single men and women of the same skill types, despite the wage rates of single women clearly having been lower.

One other purpose of this paper was to customize a coherent

[16]Boppart (2014) shows that the ratio between market services and goods increased over many decades despite an increase in the relative price of market services. In this paper, I show that this same conclusion can be drawn if market services are compared with nondurable goods. In fact, it is nondurable goods that explain the development pointed out by Boppart, as the market services-to-durable goods ratio has been monotonically declining over time.

dataset that can be used in my companion paper Almgren (2023), in which I develop a theory for understanding facts about time-use and expenditures. As a simple, but potent, motivation for doing that, I in this paper show that some of the predictions from a standard economic framework with stable homothetic preferences do not square with what is observed in data.

The companion paper investigates whether a model with stable homothetic preferences and home production potentially can generate results that align with time-use and expenditure data. Specifically, the paper focuses on low and high skilled single men and women, and how time-use and expenditures on different types of consumption groups are different in the cross-section, and how they have developed over time.

References

Aguiar, M. and Bils, M. (2015). Has consumption inequality mirrored income inequality? *American Economic Review*, 105(9):2725–56.

Aguiar, M. and Hurst, E. (2007). Measuring trends in leisure: The allocation of time over five decades. *The Quarterly Journal of Economics*, 122(3):969–1006.

Aguiar, M., Hurst, E., and Karabarbounis, L. (2012). Recent developments in the economics of time use. *Annu. Rev. Econ.*, 4(1):373–397.

Almgren, M. (2023). Home production, expenditures, and time use. Technical report, IIES.

Almgren, M., Kramer, J., and Sigurdsson, J. (2023). It runs in the family: Occupational choice and the allocation of talent.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.

Coibion, O., Gorodnichenko, Y., and Koustas, D. (2021). Consumption inequality and the frequency of purchases. *American Economic Journal: Macroeconomics*, 13(4):449–82.

Costa, D. L. (2000). From mill town to board room: The rise of women's paid labor. *Journal of Economic Perspectives*, 14(4):101–122.

Dal Bó, E., Dal Bó, P., and Snyder, J. (2009). Political dynasties. *The Review of Economic Studies*, 76(1):115–142.

Eckstein, Z. and Lifshitz, O. (2011). Dynamic female labor supply. *Econometrica*, 79(6):1675–1726.

Greenwood, J., Guner, N., Kocharkov, G., and Santos, C. (2016). Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation. *American Economic Journal: Macroeconomics*, 8(1):1–41.

Greenwood, J., Seshadri, A., and Yorukoglu, M. (2005). Engines of liberation. *The Review of Economic Studies*, 72(1):109–133.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4):1693–1754.

McGrattan, E. R., Rogerson, R., et al. (2004). Changes in hours worked, 1950-2000. *Federal Reserve Bank of Minneapolis Quarterly Review*, 28(1):14–33.

Ramey, V. (2007). How much has leisure really increased since 1965? *University of California at San Diego Working Paper*.

Ramey, V. A. (2009). Time spent in home production in the twentieth-century united states: New estimates from old data. *The Journal of Economic History*, 69(1):1–47.

Ramey, V. A. and Francis, N. (2009). A century of work and leisure. *American Economic Journal: Macroeconomics*, 1(2):189–224.

Appendices

1.A Appendix

1.A.1 Assortative Mating: A Comparison With Greenwood et al. (2016)

Greenwood et al. (2016) quantify changes in assortative mating by running the following regression (see their Equation (1))[17]:

$$\begin{aligned} \mathsf{EDU}^f_t = &\alpha + \beta \times \mathsf{EDU}^m_{t0} + \sum_{y \in \mathcal{Y}} \gamma_t \times \mathsf{EDU}^m_t \times D_{y,t} \\ &+ \sum_{y \in \mathcal{Y}} \theta_t \times D_{y,t} + \varepsilon_t \end{aligned} \tag{1.7}$$

EDU^f and EDU^m are dummy variables and take on values of one for high skilled men and women, respectively. Note that subscript $t0$ represents some base year (1962 in my case and 1960 in Greenwood et al. (2016)), and hence $\mathcal{Y}$ omits $t0$. $D_{y,t}$ is a dummy variable that takes on the value of one if $y = t$ and is zero otherwise. The constant, α, captures the share of women with low skilled men who are high skilled in the base year, and $\alpha + \beta$ captures the share of women with high skilled men who are high skilled in the same year. Coefficients θ capture time-fixed effects and in each year t the share of high skilled women in couples in which men are low skilled is given by $\alpha + \theta_t$. Compared to the base year, the difference in the share of low skilled women in couples in which men are low skilled is given by θ_t. In couples in which the man is high skilled, the corresponding difference is given by $\theta_t + \gamma_t$. The difference between these two differences, i.e.,

[17]Note that I do not follow their exact notation, but it should still be easy to compare.

γ_t, is interpreted as a measure of how assortative mating changes between $t0$ and t. They find that it increases over time; between the years 1960 and 2005 the increase is around 0.2, i.e. $\gamma_{2005} \approx 0.2$ (see their Figure 2).

I run the same regression and find similar results, which are shown in Figure 1.22. The increase that I find until the year 2005 is about 0.22, i.e., very close to their result. One thing to note is that the estimated values for the γ-vector are exactly the same as the difference between the two values for high and low skilled men, in each respective year, in Figure 1.3a.

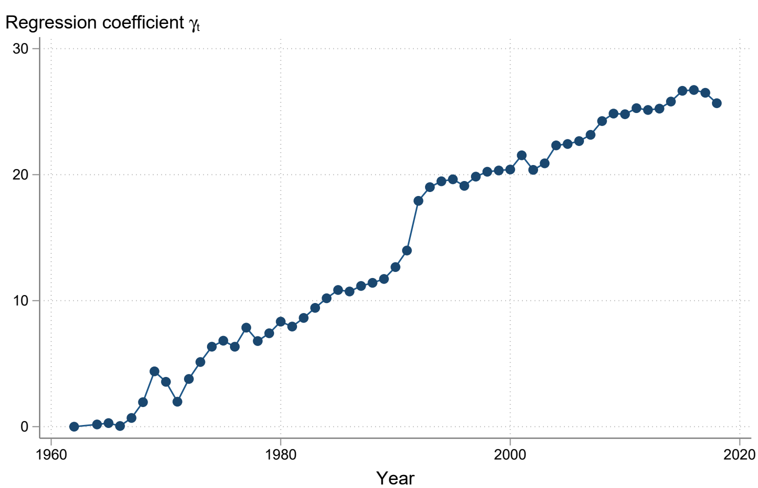

Figure 1.22: *Assortative mating like in Greenwood et al. (2016)*

Note: The values represent estimates for γ_t from Equation (1.7).

One problem with the above measure of assortative mating is that it can mechanically increase when the overall share of high skilled women in couples increases. To see this, consider the following. Denote the share of high skilled women in couples in any period t by λ_t^f. Moreover, let ϕ_t be the share of high skilled women that couple with high skilled men. The rest, $1-\phi_t$, couple with low skilled men. The share of low skilled men in couples that couple with high skilled women is then $\lambda_t^f(1-\phi_t)$, and for high skilled men the share is $\lambda_t^f\phi_t$. The difference in differences here, between some base year $t0$ and $t > t0$, which then corresponds to γ_t, is equal to

$$\gamma_t = \left[\lambda_t^f\phi_t - \lambda_{t0}^f\phi_{t0}\right] - \left[\lambda_t^f(1-\phi_t) - \lambda_{t0}^f(1-\phi_{t0})\right] \tag{1.8}$$

Assume now that the share of high skilled women that couple with high skilled men is constant, i.e., that $\phi_{t0} = \phi_t = \phi$, then it collapses to

$$\gamma_t = (\lambda_t - \lambda_{t0})(2\phi - 1) \tag{1.9}$$

γ_t will be positive if $\lambda_t > \lambda_{t0}$ (the share of high skilled women increases) and $\phi > 0.5$ (high skilled women are more likely to couple with high skilled men). That the share of high skilled women increases over time has already been established. Something that is not obvious is that $\phi > 0.5$ in all years that I consider, even at the beginning when the fraction of high skilled men is very low. E.g., in the year 1962 more than 60 percent of high skilled women couple with high skilled men. My point here is that γ can increase even if high skilled women, from a probability point of view, are not more likely to couple with high skilled men.

Now introduce λ_t^m and let it represent the share of high skilled men in period t. The measure that I propose captures to what extent sorting deviates from what would be the outcome under randomness. With random sorting, the fraction of high skilled women that couples with high skilled men would simply be λ_t^m, and the fraction that couples with low skilled men would be $1-\lambda_t^m$. The deviation from random sorting can be interpreted as a bias, which for high skilled is ϕ_t/λ_t^m and for low skilled is $(1-\phi_t)/(1-\lambda_t^m)$.[18] The bias ratio that I propose, and which is plotted in Figure 1.3b, is calculated as:

$$\text{bias ratio} = \frac{\phi_t/\lambda_t^m}{(1-\phi_t)/(1-\lambda_t^m)} \tag{1.10}$$

1.A.2 Prices

Household decisions about how to allocate their time, and their expenditures, are influenced by prices. For example, the real wage rate influences the labor supply decision: how much time does the households choose to work in the market, and how much time is left for other activities? Here, I show the relative prices between nondurable goods and services, and durable goods and services, respectively.

The data comes from the National Income and Product Accounts (NIPA).[19] The relative prices of nondurable and durable goods, respectively, fell between 1962 and 2018. Figure 1.23 shows that the fall is substantial for nondurable goods, but many times larger for durable goods. The relative price between durable goods and services was more than four times higher in 1962, compared to in 2018. For nondurable goods, the relative prices were about 1.7 times higher. Note, however, that the price index for durable goods that is presented here is reflecting the change in the price for purchasing *new* durable goods. But it is not necessarily a good measure of the price

[18]Note that this type of bias has been used in the literature before. See, e.g., Dal Bó et al. (2009) and Almgren et al. (2023).

[19]Specifically, I use the time series for the three groups: *Durable goods* (line 3), *Nondurable goods* (line 9), and *Services* (line 13) from NIPA Table 2.3.4.

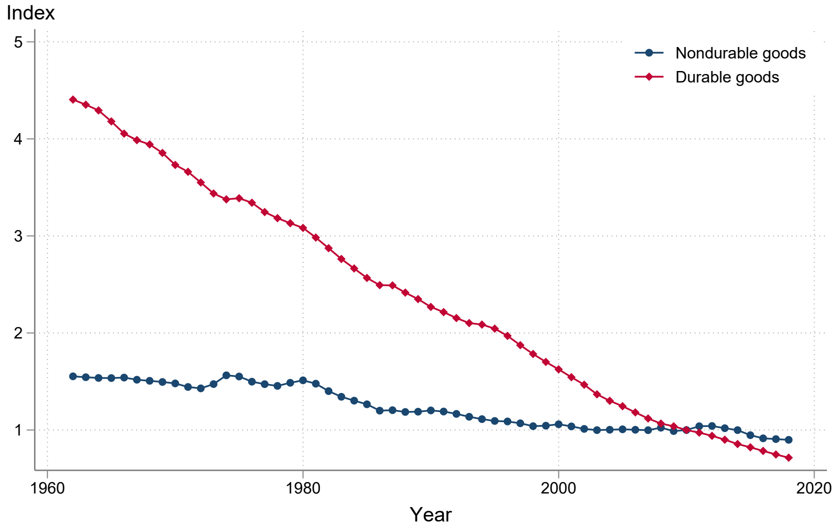

Figure 1.23: *Prices of two different types of goods, relative to services*

Note: Data comes from the NIPA tables. Indexed series with the base year 2010. See the text for more details.

of the *existing stock* of durable goods that the household owns. In Section 1.A.8, I compute the *cost of capital*.

In Section 1.2.3 it is shown that $P \equiv \left(\omega_m + P_n^{1-\sigma}\omega_n + P_k^{1-\sigma}\omega_k\right)^{\frac{1}{1-\sigma}}$ is the implicit price index for the consumption index C. I have normalized all prices, including the hourly wage rates, by the service price and therefore it does not explicitly enter this price index. It is easy to see that P is lower in the year 2018 compared to in the year 1962, since both P_n and P_k have decreased.

1.A.3 Mapping Between Gross and Disposable Labor Income

I map *labinc* (denoted by y_{pre}) to disposable income (denoted by y_{post}) by using the tax-and-transfers function from Heathcote et al. (2017): $y_{post} = \lambda y_{pre}^{1-\tau}$. The progressivity of the system is captured by τ, which I set to 0.181, the same value as the inventors of the function use. In order to calibrate λ, which I let vary by year, I use *NIPA Table 2.1 Personal Income and Its Disposition*. From this table, and for each year, I calculate an adjusted measure of disposable income. This measure excludes *Personal income receipts on assets*, *Rental income of persons with capital consumption adjustment*, and *Supplements to wages and salaries* from *Disposable personal income*.[20] Denote this measure in the NIPA table by Y^*_{post}. The sum of two rows in the table: *Wages and salaries* (row 3), and *Proprietors' income with inventory valuation and capital consumption adjustments* (row 9) should map closely to the variable *labinc* that I calculate using the CPS data. Denote this by Y^*_{pre}. Now I can calculate the share Y^*_{post}/Y^*_{pre}, which, for each year, will serve as the targeted statistic for λ. Specifically, I find the value of λ that solves the following equation:

$$\frac{\lambda \sum_i y_{pre,i}^{1-\tau}}{\sum_i y_{pre,i}} = \frac{Y^*_{post}}{Y^*_{pre}} \tag{1.11}$$

where I sum over all the individuals $i \in I$ in the CPS data, in each respective year.

[20]Referred to by their line numbers, I exclude lines 6, 12, and 13 from line 27.

1.A.4 Women's Skill Premium With Age Fixed Effects

In Section 1.2.4, I look at the difference in the earnings of high skilled single women and low skilled single women, and refer to this difference as the "skill premium". However, Figure 1.14a and Figure B.6a are based on simple averages: one for each year and skill level, while no other differences between the two groups are considered. A possible explanation for the drop in the difference between wage rates of high skilled and low skilled single women that is observed in the 1970s could, e.g., be that the average age of high skilled fell relative to the average age of low skilled. To control for such heterogeneity, I run the following regression:

$$\begin{aligned} y_{it} &= \alpha + \beta_a \times age_{it} + \gamma_t \times year_{it} + \delta * high_i \\ &\quad + \eta_t \times (high_i \times year_{it}) + \varepsilon_{it} \end{aligned} \tag{1.12}$$

where the left-hand side variable y_{it} is the natural logarithm of the hourly wage for individual i in year t. I estimate separate βs for each age, which is highlighted by the subscript a on β, where $a = \{25, \ldots, 64\}$ depending on the individual's age. Similarly for years, I estimate separate coefficients for each year $t = \{1963, \ldots, 2018\}$. The variable $high$ is a dummy variable that takes the value of one if the individual is high skilled, and otherwise takes the value of zero. Hence, δ is a measure of the skill premium in year 1962. Lastly, η_t with $t = \{1963, \ldots, 2018\}$ is a set of year-specific coefficients that capture the change in the skill premium since the year 1962.

Using the results from this regression, I then predict the average wages of individuals in each respective skill group after imposing that they are all of age 40 in each year, and thus eliminating the differences in wage rates that are driven by differences in age. I then transform it back to real wages and compute indexed series. These are shown in Figure 1.24. Although they are noticeably different from the series in B.6a, they are qualitatively very similar. I do conclude that heterogeneity in how age composition develops in the two groups did indeed contribute to differences in wage developments between the two skill groups, but that it is not the major driver.

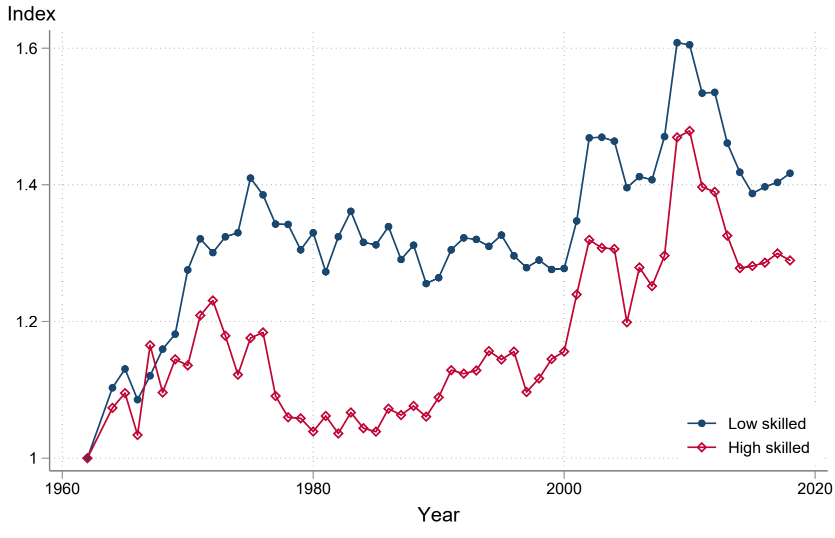

Figure 1.24: *Growth in hourly wage rates, with age fixed effects*

Note: These plot the developments of hourly wages rates for low and high skilled single women after the effects from age composition have been removed. Both series are indexed to their values in the year 1962. See the text for more details.

1.A.5 Effects From Age Composition on Wage Rates Among Men in Couple and Single Households

I do the same type of exercise as in Appendix 1.A.4, but instead of comparing groups with different skill levels, I now instead compare single with couple households. For completeness, I write out the full specification that I run in order to isolate the developments in earnings that are *not* driven by changes in age composition. I estimate the coefficients in the following regression:

$$\begin{aligned} y_{it} &= \alpha + \beta_a * age_{it} + \gamma_t * year_{it} + \delta * couple_i \\ &\quad + \eta_t * (couple_i \times year_{it}) + \varepsilon_{it} \end{aligned} \tag{1.13}$$

I will not comment on the parts of this specification that are identical to Equation 1.12, for this I refer the reader to the text commenting on that equation.

$couple$ is a dummy variable and takes on the value of one if the man lives in a couple household. δ captures how much higher wages that men in couple households have compared to men in single households, eliminating the differences in the age composition, in the year 1962. The coefficients η_t with $t = \{1963, \ldots, 2018\}$ are estimates of how this difference then developed until the year 2018. As in Section 1.A.4, I hold constant the age dimension by predicting earnings for a hypothetical 40-year old in each group. Figures 1.25a and 1.25b illustrate the results.

The difference in earnings between men in couple households and men in single households is similar. However, we can learn that in many years differences in the age composition have a noticeable effect. E.g., while the comparison of means yields a relative difference of about 4 percent in 1962 and then 10 percent in 2018, the relative differences were about the same at 10 percent in both of these years when holding age constant.

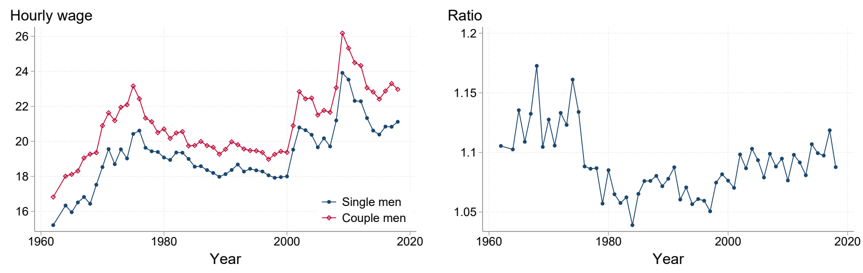

(a) *Wage rates, by status* ***(b)*** *Ratio*

Figure 1.25: *Removing how age affects average wage rates*

Note: *Panel (a)*: Shows fitted average wages, by year, for men in single households and couple households, respectively. The cross-sectional differences and developments over time are cleansed from effects due to age composition. See the text for details. *Panel (b)*: Indexed versions of the series in panel (a).

1.A.6 Effects From Age Composition on Wage Rates Among Men in Different Types of Households

I estimate yet another regression to control for heterogeneity in age composition. The regression estimated it:

$$
\begin{aligned}
y_{it} = \alpha + \beta_a \times age_{it} + \gamma_t \times year_{it} + \delta \times type_i \\
+ \eta_t \times (type_i \times year_{it}) + \varepsilon_{it} \qquad (1.14)
\end{aligned}
$$

where the coefficients of interest are the δs and the ηs. Different from in earlier sections, δ is now a vector with five coefficients, one for high skilled single men and the remaining four for the four different types of couple households. Low skilled single men are absorbed by the constant. η contains a large number of coefficients, one for each combination of the five household types (the time effect for low skilled single men is absorbed by the vector γ) and year (except for the first year, that is absorbed by the constant). I again produce time series for a hypothetical 40-year old in each respective type of household, and the results are shown in Figure 1.26.

Some differences compared with the unadjusted counterpart (Figure B.2a) are clearly visible. The most noticeable is that the skill premium for single men was much higher in the year 2018 compared to in the year 1962 when just comparing means, but if the differences in age composition are removed this is no longer the case. Among couple households there are clearly differences, but overall the picture is similar to when just comparing means.

Comparison With Results on Average Hours Worked in Home Production In Ramey (2009), and Aguiar and Hurst (2007)

In Figure 1.27 I compute average hours worked in home production among men and women, respectively, and compare them with the results in Ramey (2009), and Aguiar and Hurst (2007). Starting with Aguiar and Hurst, they are similar, as they should be since the data source is the same and I follow their scheme when classifying activities. Although the differences are small, I will briefly comment on their potential underlying reasons. First, the sample is somewhat different. While they consider individuals who are aged between 21-65, the range that I consider is 25-64. Second, they adjust for demographic changes while I do not. Clearly, these adjustments make little difference, however.

The differences as compared to the results in Ramey (2009) are

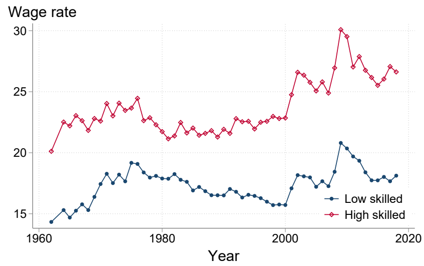

(a) Single men

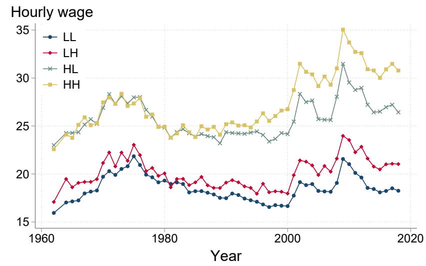

(b) Men in couples

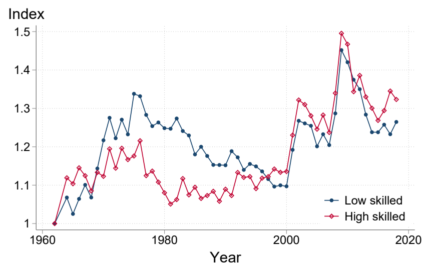

(c) Single men, indexed

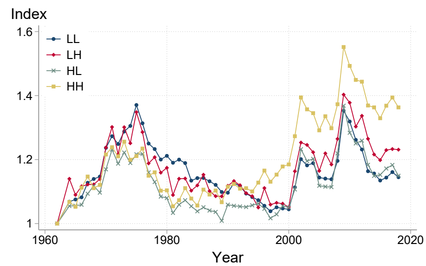

(d) Men in couples, indexed

Figure 1.26: *Wage rates for men in different types of households, with age fixed effects*

Note: *Panel (a)*: Average wage rates among two groups of single men: low skilled (blue circles) and high skilled (red diamonds). *Panel (b)*: Average wage rates among men in four different kinds of couple households. The legend in the graph tells the combination of the man's as well as the woman's skill level. The first letter indicates the skill level of the man and the latter the skill level of the woman; e.g., LH (red diamonds) represents the development of average hours worked in the market for a man who is low skilled and whose partner is high skilled. *Panels (c) and (d)*: indexed series of the lines in panels (a) and (b), respectively.

larger. In terms of changes over time, it is especially after the year 1985 that Ramey finds different results; while I, and Aguiar and Hurst, find that the hours worked in home production decreased for women and showed no clear trend for men, Ramey finds that they increased for men, and increased slightly for women. The differences between Ramey (2009) and Aguiar and Hurst (2007) are well known and originate from how they differentially categorize activities. Ramey (2007) criticizes some of the choices made by Aguiar and Hurst (2007), and Aguiar and Hurst later respond to this critique. I will not enter this debate or give a description of the details, but merely, instead, acknowledge that the trends in time spent on different activities might be sensitive to the categorization.

1.A.7 Effects From Age Composition on Hours in Home Production

Time spent in home production might be correlated with age. I control for differences in the age composition, and differential trends in it, across groups by running the following regression:

$$\begin{aligned} y_{it} &= \alpha + \beta_a * age_{it} + \gamma_t * year_{it} + \delta * type_i \\ &\quad + \eta_t * (type_i \times year_{it}) + \varepsilon_{it} \end{aligned} \tag{1.15}$$

where y_{it} represents hours in home production of individual i in year t and $type_i$ is an indicator for i's household type, being one out of the four possible skill level (low/high) $\times$ relationship status (single/couple) combinations. With the results from this regression, I predict hours in home production for each individual but hold age constant by setting it to 40 for everyone. The resulting cross-sectional differences and the trends (and differences in trends) are then not a result of the heterogeneity in age composition.

The results are shown in Figure 1.28. For women, the results are very similar, in all different dimensions, compared to the unadjusted averages. However, they look different for men. The difference that

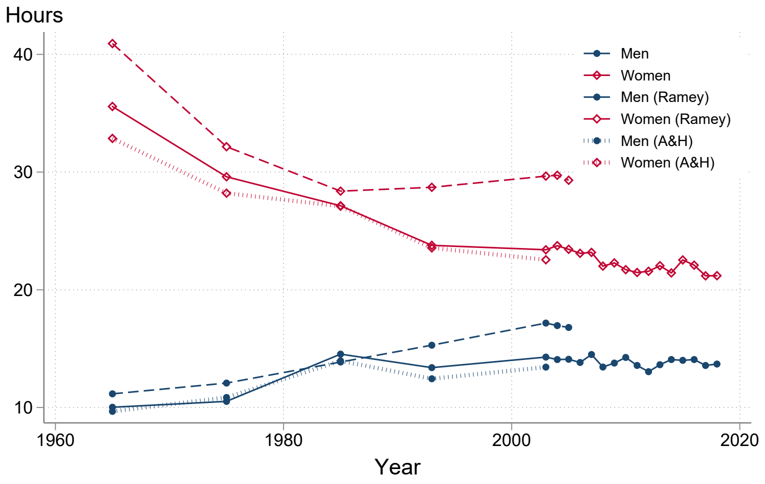

Figure 1.27: *Comparison with Ramey (2009), and Aguiar and Hurst (2007)*

Note: Different estimates of average hours worked in home production among men and women, respectively. The markers represent the point estimates, and the lines simply represent linear interpolations between these points. Importantly, note that Ramey in her paper actually estimates values for years in between these points, and that these estimates result in values which are not identical to the one suggested by the linear interpolations shown here. Markers that are connected with solid lines represent my estimates, markers connected with dashed lines represent estimates by Ramey, and markers connected with dotted lines represent estimates from Aguiar and Hurst.

stands out is that the increase in hours between the early 1960s and 2010s is around twice as big for single men. For men in couple households, the increase is around 50 percent larger when changes in age-composition are taken into account.

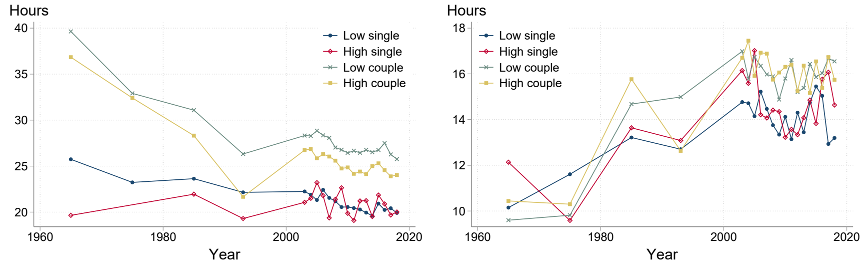

(a) Women *(b) Men*

Figure 1.28: *Hours Worked in Home Production, With Age Fixed Effects*

Note: *Panel (a)*: Hours worked in home production for women, depending on their skill levels and whether they live in a couple or single household. *Panel (b)*: Hours worked in home production for men, depending on their skill levels and whether they live in a couple or single household. Age effects are taken out. See the text for details.

1.A.8 Expenditure Shares

For the years 1990-2018, I download the CEX data directly from the U.S. Bureau of Labor Statistics.[21] I use the mtbi files for expenditures, and the family characteristics and income files to gather information about education and age of the head as well as a potential spouse. For the years 1980-1989, I download the corresponding expenditure and family characteristics files from the ICPSR.[22]

Categorizing Expenditures in the CEX

I exclude information from the mtbi files in which I have no interest, such as information about assets. To identify which rows to drop from mtbi files, I make use of the *Hierchical groupings* that are made

[21] www.bls.gov/cex/
[22] https://www.icpsr.umich.edu/web/ICPSR/series/20

available by the BLS.[23] These files exist from the year 1997. For the earlier years, I use the hierchies from the year 1997. These hierchical groupings do not change by much over time, so this should work well. I exclude all rows from the mtbi files that are not classied as "FOOD" (food expenditures) or "EXPEND" (non-food expenditures).

How large a part of expenditures are allocated to nondurable goods, durable goods, and services, respectively, for different types of households? I rely on the already existing work by Coibion et al. (2021) to answer this question. Their Online Appendix B includes a comprehensive list of UCC codes that appear in the Interview Surveys and, importantly, they assign UCCs to expenditure categories. To a large extent, I follow Coibion et al. (2021) when assigning UCCs to the major spending categories.

UCCs are divided into subgroups and also into main spending categories. One example of a subgroup is "Household appliances" which, in turn, consists of 28 unique UCCs. This is, in turn, one of 30 subgroups that all consist of UCCs that are classified as durable goods. Because the full list of UCCs is long, I do not include it here. Instead, I will specify the additions and changes that I make to the list from Coibion et al. (2021). In Table 1.2, I list all the UCCs that appear in the mtbi (expenditure files) but were not included in Coibion et al. (2021), in most cases because their analysis ends in the year 2016 and the UCC was added after that. Column 2 in the table gives information about into which main expenditure category that I sort the UCC. To sort, I use the data dictionary[24] for the CEX, which gives a short description of each code value and is specified in column 5.

For about ten percent of the UCCs, I sort the UCC into a main

[23]https://www.bls.gov/cex/pumd_doc.htm

[24]https://www.bls.gov/cex/pumd/ce_pumd_interview_diary_dictionary.xlsx

Table 1.2: *Added UCCs*

(1) UCC	(2) Main group	(3) Description
550320	Durable	Medical equipment for general use[1]
550340	Durable	Hearing aids[1]
190904	Nondurable	Food prepared by consumer unit on out-of-town trips[1]
410110	Nondurable	Infant coat, jacket, snowsuit[1]
410120	Nondurable	Infant dresses, outerwear[1]
410130	Nondurable	Infant underwear[1]
410120	Nondurable	Infant nightwear, loungewear[1]
470311	Nondurable	Electric vehicle charging[2]
480216	Service	Vehicle clean services including carwash[1]
620215	Service	Ticket to movies[2]
620216	Service	Tickets to parks or museums[2]
310243	Service	Rental, streaming, downloading video[2]
560410	Service	Non physician services inside home[2]
560420	Service	Non physician services outside home[2]
580401	Service	Long term care insurance (not BCBS)[2]
580402	Service	Long term care insurance (BCBS)[2]
580411	Service	Dental care insurance (not BCBS)[2]
580412	Service	Dental care insurance (BCBS)[2]
580421	Service	Prescription drug insurance (not BCBS)[2]
580422	Service	Prescription drug insurance (BCBS)[2]
580431	Service	Vision care insurance (not BCBS)[2]
580432	Service	Vision care insurance (BCBS)[2]
580441	Service	Vision care insurance (not BCBS)[2]
580442	Service	Vision care insurance (BCBS)[2]
580906	Service	Other health insurance (BCBS)[2]
580908	Service	Medicaid premiums
580909	Service	Tricare/military premiums
580910	Service	Tricare/military premiums
002120	Service	Other non-health insurance[1]

Note: *Column (1)*: Added UCCs; *column (2)*: main expenditure category into which I sort the UCC; *column (3)*: desription of UCC, taken from the CEX data dictionary. Superscript 1 indicates that the UCC was not listed in Appendix Table B1 in Coibion et al. (2021) but did exist as a UCC. Superscript 2 indicates that the UCC was added in year 2017.

expenditure category that is different as compared to how Coibion et al. (2021) sort it. I list these UCCs in Table 1.3 and specify into which category that they sorted the UCC, and how I do it, in columns 2 and 3, respectively. As mortgage payments are a form of savings, I categorize them as *Nonconsumption*, whereas they grouped them as a service expense. For the rest of the UCCs, I choose to move UCCs that they listed as Nonconsumption to the Service group instead.

Table 1.3: *Different main expenditure categories*

(1) UCC	(2) From	(3) To	(4) Description
830101	Service	Noncons	Special lump sum mortg. payment (owned home)
830102	Service	Noncons	Special lump sum mortg. payment (owned vacation)
790910	Service	Noncons	Special lump sum mortg. payments (other property)
220111	Noncons	Service	Fire and extended coverage
220112	Noncons	Service	Fire and extended coverage
220121	Noncons	Service	Homeowners insurance
220122	Noncons	Service	Homeowners insurance
350110	Noncons	Service	Tenant”s insurance
500110	Noncons	Service	Vehicle insurance
580110	Noncons	Service	COMMERCIAL HEALTH INSURANCE
580210	Noncons	Service	BLUECROSS/BLUE SHIELD
580310	Noncons	Service	Health maintenance plans
580901	Noncons	Service	Medicare payments
580902	Noncons	Service	COML MEDICAR SUPLMNT/OTH HLTH INS
700110	Noncons	Service	Life, endowment, annuity, other personal insurance
580111	Noncons	Service	Traditional fee for service health plan (not BCBS)
580112	Noncons	Service	Traditional fee for service health plan (BCBS)
580113	Noncons	Service	Preferred provider health plan (not BCBS)
580114	Noncons	Service	Preferred provider health plan (BCBS)
580311	Noncons	Service	Health maintenance organization (not BCBS)
580312	Noncons	Service	Health maintenance organization (BCBS)

Continued on next page

Table 1.3 – *Continued from previous page*

(1) UCC	(2) From	(3) To	(4) Description
580903	Noncons	Service	Commercial medicare supplement (not BCBS)
580904	Noncons	Service	Commercial medicare supplement (BCBS)
580905	Noncons	Service	Other health insurance (not BCBS)
585906	Noncons	Service	Can not find in data dictionary
580400	Noncons	Service	Long term care insurance
580907	Noncons	Service	Medicare prescription drug premium
580115	Noncons	Service	Fee for service health plan (not BCBS)
580116	Noncons	Service	Fee for service health plan (BCBS)
220210	Noncons	Service	Can not find in data dictionary
220211	Noncons	Service	Property taxes
220212	Noncons	Service	Property taxes
950024	Noncons	Service	VEHICLE PERSONAL PROPERTY TAXES
520110	Noncons	Service	VEHICLE REGISTRATION STATE/LOCAL
520310	Noncons	Service	Drivers" license
520410	Noncons	Service	Vehicle inspection
520901	Noncons	Service	Docking and landing fees
620110	Noncons	Service	CLUB MEMBERSHIP DUES AND FEES
620121	Noncons	Service	Fees for participant sports
790630	Noncons	Service	Special assessments (other property)
840101	Noncons	Service	Special assessments (owned home)
840102	Noncons	Service	Special assessments (owned vacation)
450311	Noncons	Service	Charges other than basic lease, such as insurance or maintenance (car lease)
450411	Noncons	Service	Charges other than basic lease, such as insurance or maintenance (truck/van lease)
620111	Noncons	Service	Social, recreation, health club membership
620112	Noncons	Service	Credit card memberships
620113	Noncons	Service	Automobile service clubs
620114	Noncons	Service	Automobile service clubs and GPS services
220311	Noncons	Service	Mortgage interest
220312	Noncons	Service	Mortgage interest
510110	Noncons	Service	Automobile finance charges

Continued on next page

Table 1.3 – *Continued from previous page*

(1) UCC	(2) From	(3) To	(4) Description
510901	Noncons	Service	Truck finance charges
510902	Noncons	Service	Motorcycle and plane finance charges
680220	Noncons	Service	Checking accounts, other bank service charges
710110	Noncons	Service	Finance charges excluding mortgage and vehicle
850300	Noncons	Service	Other vehicle finance charges
220313	Noncons	Service	Interest paid, home eq. loan
220314	Noncons	Service	Interest paid, home eq. loan
880110	Noncons	Service	Interest paid, home eq. line of credit
880210	Noncons	Service	Interest paid, home eq. line of credit (other prprty)
880310	Noncons	Service	Interest paid, home eq. line of credit
220321	Noncons	Service	Prepayment penalty charges
220322	Noncons	Service	Can not find in data dictionary

Note: *column (1)*: UCC; *column (2)*: main expenditure category that Coibion et al. (2021) sort the UCC in to; *column (3)*: main expenditure category that I sort the UCC in to; *column (4)*: desription of UCC, taken from the CEX data dictionary.

I use the "rental equivalence" amount as a housing expenditure for homeowners, and include it as a service expenditure.[25] This information is not available in the mtbi files, but instead in the fmli files. Including this value (i.e., the rental equivalence amount) means that I should drop other UCCs that should implicitly be captured by the rental equivalence, e.g., mortgage interest payments. In the *Hierchical groupings*, there is one category "OWNDWELL" that consists of subgroups that all are related to the costs of owning and maintaining a house. I exclude all UCCs in this category. However, before

[25]The household is answering the following question: "If someone were to rent your home today, how much do you think it would rent for monthly, unfurnished and without utilities?"

I remove them, I create an indicator variable that is equal to one if the household has positive spending on any of the UCCs in the OWNDWELL category. There are some households for which there is no information about the rental equivalence amount, even though they report positive values in the OWNDWELL cateogory and thus most likely are homeowners. Furthermore, before the year 1985 this question did not exist. I use information about other types of expenditures and impute the rental equivalence amount. More specifically, I first regress the rental equivalence amount on: (i) nondurable expenditures, (ii) durable expenditures, (iii) service expenditures, and (iv) year. All variables are in logs. I then predict the rental equivalence amount for households for which this information is missing, but have reported positive spending on the OWNDWELL category. The R2 from this regression is 0.51. Lastly, I transform the predicted variable by taking the exponential.

CEX Expenditure Shares for Different Demographic Groups

Now when expenses have been grouped into categories, it is possible to, for each household, calculate the fractions of expenditures that are allocated to each of these categories: nondurable goods, durable goods, and services. By using information from the fmli files, I attach a skill, low or high, to individuals. I then separate households into groups, based on whether it is a single household or a couple household, and depending on skill levels. Households are also grouped by age: either placed as a working-age household (ages 25-64), or as a household in retirement age (65 or older). My focus in this paper is on working age households. However, I do later make adjustments to the expenditure shares in order to match aggregate statistics from the national accounts. In short: I acknowledge that it would be inappropriate to directly match the aggregate statistics, as it is influenced by the levels and the allocation of spending among the retired group. I will go into some more detail about this in Section 1.A.8.

Before collapsing the data at the household type times year level, I drop observations that do not report strictly positive expenditure shares on all three expenditure categories. Then, the expenditure shares in each type of household that I use is the mean value. The means of the expenditure shares do not sum to one and therefore one final adjustment is made to ensure that they do. If η_j^1 is the expenditure share on category j, where j is a generic index, before the adjustment, then the adjustment is simply $\eta_k = \frac{\eta_k^1}{\sum_j \eta_j^1}$. For the group of households that consists of one low skilled man and one high skilled woman, the mean durable goods expenditure share is particularly low in year the 1980. The expenditure share in this group deviates in a striking way compared to how it typically moves in relation to the expenditure shares in other household types. I cannot find any explanation for this, but nevertheless choose to consider it an outlier and replace the expenditure shares for this particular household type with their values from the year 1981.

The CEX data starts in the year 1980, and hence, without any further assumptions, I cannot produce any household specific time series before that point. I make the assumption that, across household types, expenditures on each expenditure category are as in year 1980. Without any further adjustments, this of course means constant expenditure shares. However, in Section 1.A.8, I will present additional adjustments that cause these expenditure shares to vary also in the 1962-1979 period. Next, I will go into detail about how I adjust them to (i) make them line up with aggregate data from the National Accounts, and (ii) take into account the user cost of capital rather than expenses on newly purchased durable goods.

Adjustments to Match Aggregate Statistics

A discrepancy has been documented between aggregate expenses in the NIPA and the population weighted expenses in the CE Survey. A second concern is that some types of consumption seem to be under-reported more than others, and that this problem has increased over time (see, e.g., Aguiar and Bils (2015)). To decrease the issues of under-reporting, I adjust the expenditure shares, for each skill type in each year, such that the type-weighted averages match the aggregate statistics.

The first step involves combining the shares with data from the National Accounts. For the aggregate statistics, I use NIPA Table 2.3.5. "Personal Consumption Expenditures by Major Type of Product", from which I read off aggregate expenditures on nondurable goods, durable goods, and services, respectively. I calculate an adjusted per-capita expenditure statistic. It is a rough approximation but it should be superior to the unadjusted counterpart. I want to get a value for per-capita expenditures for individuals in the types of households on which I focus. To this end, I need to know two things: (i) the fraction of individuals in the age group 25-64, and how much they, on average, spend relative to individuals in the age group 65+. I do implicitly make the assumption that the contribution made by individuals younger than 25 years old, to aggregate expenditures, is negligible. This is surely at odds with the data. However, at this point the assumption will not have any strong implications for the results, as will later become clear. Let E_j denote aggregate expenditures on the main expenditure group $j \in \{n, k, m\}$, and let ψ^{young} and ψ^{old} denote the population shares of individuals in age groups 25-64, and

65+, respectively. Furthermore, let $\tilde{E}^{young}$ and $\tilde{E}^{old}$ denote the average expenditures for each age group in the CEX. I suppress time indicators, but note that all these statistics are year-specific. Now, the adjusted per capita expenditure for individuals in the age group 25-64 on main expenditure group j is given by

$$\bar{E}_j = \frac{\psi^{young}\tilde{E}^{young}}{\psi^{young}\tilde{E}^{young} + \psi^{old}\tilde{E}^{old}} \times \frac{E_j}{Q^{young}} \tag{1.16}$$

where Q^{young} is the number of individuals in the 25-64 age group. Hence, the first part on the right-hand-side of Equation (1.16) represents the share of expenditures coming from the young age group.[26] By multiplying that by E_j, I get an approximation for total expenditures on type j in the young age group, and, finally, by dividing through by Q^{young} I arrive at the desired per capita statistic.

The next step involves adjusting the household type-specific expenditures such that they line up with the (adjusted) per-capita expenditures from the NIPA table. Let ψ_h denote the population share of individuals in household type h. Note that this is a measure of the share of *individuals* in each type of household. The reason why this is the appropriate measure is because I measure expenditures as per-capita expenditures, simply by dividing total expenditures by two in couple households. Let the per-capita expenditure on main expenditure group, that I have from the CEX data, for individuals in household type h be $\tilde{E}_{hj}$. For each group j, I adjust the levels such

[26]This has fallen considerably over the time period. While it was around 85% in the year 1980, it has fallen to around 75% in the year 2018. This is partly explained by the relative size of the 65+ age group having increased, but a big part of it is also explained by the per capita consumption in the 65+ age group, relative to the per capita consumption in the 25-64 age group, having increased significantly.

that

$$\bar{E}_j = \sum_h \psi_h \bar{E}_{hj} \tag{1.17}$$

which is simply done by setting the per capita expenditures for each individual in household type j to

$$\bar{E}_{hj} = \frac{\bar{E}_j}{\sum_h \psi_h \tilde{E}_{hj}} \tag{1.18}$$

This ensures that the expenditure shares on each main expenditure group match the expenditure shares in the adjusted aggregate, while keeping intact the relative differences in category-specific expenditures that were observed in the CEX.[27] These are not the final expenditure shares, however, as one important step remains.

Imputing Capital Stocks and User Costs of Capital

So far, I combined the CE Survey with the NIPA to impute households' expenses on durable goods. These expenses are capturing the amount that households spend on purchasing *new* durable goods, but are not necessarily a good measure of (i) the quantity of capital that the household uses or (ii) the cost of that capital. For these two reasons, I impute (i) the capital stock, for each household type and year, and (ii) the cost per unit of capital.

For what will follow, I impute time series for capital by using durable expenditures from Section 1.A.8. The CE Survey is a rotating panel survey. For this to be a valid procedure, I assume that the measured expenditures on durable goods in any period t by house-

[27]E.g., if household a spends two times the dollar amount on nondurable goods compared to household b, then this will also be the case after the adjustment.

holds of type s are similar to the expenses of households (also in period t) of the same type, that were interviewed in year $t+1$ but not in t.

Say that the capital stock that the household of type s has in period $t-1$ is K_{st-1}. There are three more factors that affect what its *nominal* value will be in period t: (i) the price, (ii) depreciation, and (iii) (nominal) new investments. This can be summarized by the following equation:

$$\tilde{K}_{st} = \tilde{P}_{kt}(1-\delta_t)K_{st-1} + (1-\frac{1}{2}\delta_t)\tilde{K}_{st}^{new} \tag{1.19}$$

where, to be consistent with how data is constructed in the NIPA, it is assumed that investments on average materialize in the middle of the period and hence that half of them are subject to depreciation. I use expenditures on durable goods from Section 1.A.8 for new investments, $\tilde{K}_{st}^{new}$. The first part on the right-hand-side is the undepreciated part of the capital stock in period $t-1$, at year t's price. The equation can also be written as follows:

$$\tilde{K}_{st} = \frac{\tilde{P}_{kt}}{\tilde{P}_{kt-1}}(1-\delta_t)\tilde{K}_{st-1} + (1-\frac{1}{2}\delta_t)\tilde{K}_{st}^{new} \tag{1.20}$$

where the right-hand-side now instead contains the nominal value of the capital stock in period $t-1$ and re-valuates it to express it in the next period's prices. To impute the time series for capital stocks, I need information about all the variables and prices on the right-hand-side. The value of new investments has already been computed. However, the ratio $\frac{\tilde{P}_{kt}}{\tilde{P}_{kt-1}}$ is not explicitly reported in the NIPA tables. I make use of two existing series for the capital stock: the *Current-Cost Net Stock of Consumer Durable Goods* (NIPA Table 8.1) and

Chain-Type Quanitity Indexes for Net Stock of Consumer Durables (NIPA Table 8.2). In short: I have information about both nominal and real stocks. Now I make use of the identity $\tilde{K}_t = \tilde{P}_{kt} K_t$ to find

$$\frac{\tilde{K}_t}{\tilde{K}_{t-1}} = \frac{\tilde{P}_{kt} K_t}{\tilde{P}_{kt-1} K_{t-1}} \tag{1.21}$$

$$\implies \frac{\tilde{P}_{kt}}{\tilde{P}_{kt-1}} = \frac{\tilde{K}_t / \tilde{K}_{t-1}}{K_t / K_{t-1}} \tag{1.22}$$

Return back to Equation (1.20) and solve for δ, which is now the only remaining unknown:[28]:

$$\delta_t = \frac{\frac{\tilde{P}_{kt}}{P_{kt-1}} \tilde{K}_{jt-1} + \tilde{K}^{new}_{jt} - \tilde{K}_{jt}}{\frac{\tilde{P}_{kt}}{P_{kt-1}} \tilde{K}_{jt-1} + \frac{1}{2} \tilde{K}^{new}_{jt}} \tag{1.23}$$

It is necessary to have starting values for the capital stocks, for all types of households. I turn to aggregate data and conclude that the capital stock in the year 1961 was 4.1 times larger than new investments in the year 1962. I will assume that the same relationship holds for all types of households. Now I can impute values for their capital stocks in the year 1961 and then, by applying Equation (1.20), generate time series.

Figure 1.29 shows the developments of capital stocks for each household type, and compares it with the developments of durable

[28]Note that the numerator on the right-hand-side of Equation (1.23) equals the depreciation in nominal terms. It is easy to see this by reshuffling the equation so that: $\delta \frac{\tilde{P}_{kt}}{P_{kt-1}} \tilde{K}_{jt-1} + \delta \frac{1}{2} \tilde{K}^{new}_{jt} = \frac{\tilde{P}_{kt}}{P_{kt-1}} \tilde{K}_{jt-1} + \tilde{K}^{new}_{jt} - \tilde{K}_{jt}$, where the first part on the left-hand-side of the equation represents the nominal depreciation of the part of the capital stock that the household entered the period with, and the second part represents the nominal depreciation of new investments. Hence I can use NIPA table 8.4. *Current-Cost Depreciation of Consumer Durable Goods* in the numerator when calculating the depreciation rate in Equation (1.23). During the period, the depreciation rate has varied between 17-19%, with an average of 18%.

expenditures. Expenditures are normalized by using the level of expenditures among low-skilled single men in the year 1962, and hence the lines are also informative about the level differences between the types of households. The indices for the capital stocks are normalized in the same fashion, by letting the base be the capital stock among low-skilled single men in the year 1962. Indices for expenditures and capital stocks follow similar trends over the period. By construction, this is to be expected if there are no significant trends in the new investments-to-capital ratio. The lines are noticeably different, however: the stocks are significantly less volatile than the expenditures and the differences between the two lines in certain years can be large.

The price index for durable goods in the NIPA is not suitable for capturing the cost of capital. My method for estimating the user cost of capital is based on the following idea: the household enters the period and has the possibility to issue debt to finance buying capital. The debt is bought by a foreign investor and the household needs to repay the debt, plus interest, at the end of the period. Capital has a nominal price $\tilde{P}_{kt}$ at the beginning of the period and depreciates at the rate δ. The interest rate is i_t. At the end of the period, the household sells the undepreciated part of the capital stock at the price $\tilde{P}_{kt+1}$ and spends the remaining resources on services and nondurable goods. Hence, the nominal cost of the capital stock can be expressed as

$$\mathrm{cost}_K = i_t \tilde{P}_{kt} K_t + \tilde{P}_{kt} K_t - (1 - \delta_t) K_t \tilde{P}_{kt+1} \qquad (1.24)$$

Divide through by $\tilde{P}_{kt} K_t$ on both sides and get the user cost of capital,

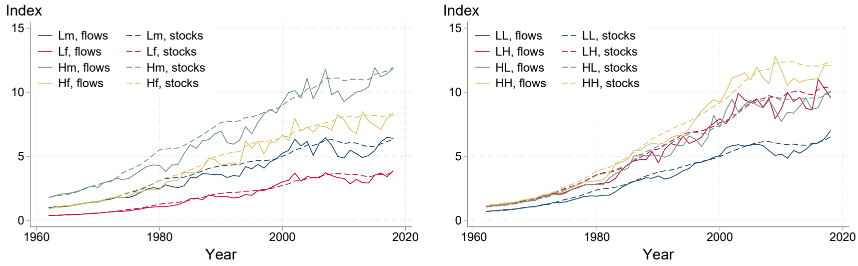

(a) *Women* ***(b)*** *Men*

Figure 1.29: *Durable goods expenditures and imputed capital stocks*

Note: The graph compares developments of imputed expenditures on durable goods (solid) and the development of imputed capital stocks (dashed) between the years 1962 to 2018. The series for single households are shown in *panel (a)*, and the series for couple households can be seen in *panel (b)*. Note that the data for couples is per capita, so the values would need to be multiplied by two to get flows and stocks at the household level. For detailed information about the construction of the series, see the text. Flows (expenditures) are normalized by using the level of flows for low skilled single men in the year 1962, and hence the lines are also informative about the level differences between the types of households. The indices for the capital stocks are normalized in the same fashion, by letting the base be the capital stock among low skilled men in the year 1962.

denoted by V:

$$V = i_t + \delta_t \frac{\tilde{P}_{kt+1}}{\tilde{P}_{kt}} - \Delta\tilde{P}_{kt} \tag{1.25}$$

where $\Delta\tilde{P}_{kt}$ denotes the percentage change in the nominal price of capital. To get the price per unit of capital, expressed in nominal currency, I multiply it by the price index $\tilde{P}_k$:

$$\widetilde{UC} = \tilde{P}_k V \tag{1.26}$$

All prices are normalized by the price of market services, $\tilde{P}_m$, and hence, to make it model consistent, I let the price per unit of capital that enters my model, P_k, be

$$P_k = \frac{\widetilde{UC}}{\tilde{P}_m} \tag{1.27}$$

I earlier adjusted expenditures to make sure that they were the same as income. By changing how capital and capital expenditures are measured, this is no longer the case. Thus I need to adjust the level again and this time the expenditure shares will change. First, for each type of household, I let the expenditures on capital be $\bar{E}_k = P_k K$, where P_k is given by Equation (1.27) and K is the index for the real capital stock. To make sure that total expenditures equal income, I will make adjustments to service and nondurable goods consumption. If excluding expenses on durable goods, each household's expenditure share on services was $\eta_{m,-k} = \frac{E_m}{E_m + E_n}$, where E_m are service expenditures and E_n are expenditures on nondurable goods. I will make sure that the same relationship holds even after the adjustment. Denote the disposable income of a household type as Y_{post}. Service expenditures will be adjusted to

$$\bar{E}_{m,s} = \eta_{m,-k}(Y_{post} - \bar{E}_k) \tag{1.28}$$

and I make sure that expenditures equal income if

$$\bar{E}_n = (1 - \eta_{m,-k})(Y_{post} - \bar{E}_k) \tag{1.29}$$

1.B Figure Appendix

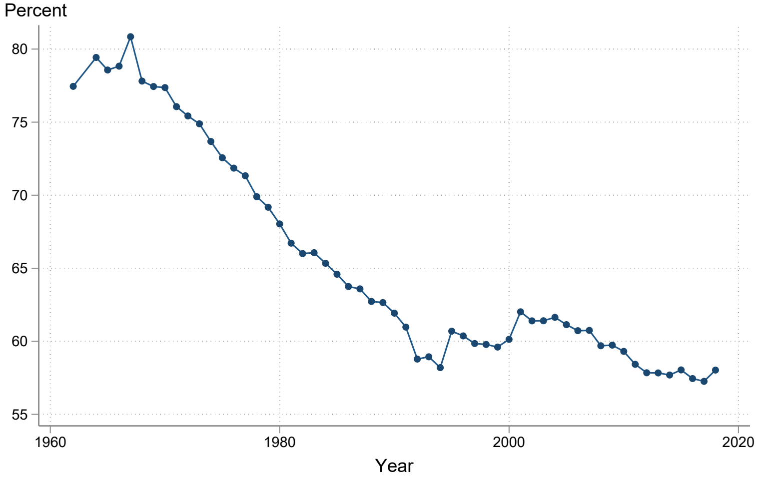

Figure B.1: *Share of individuals that live in couple households*

Note: The sample considered here consists of all individuals in the CPS between ages 25-64 for whom I observe information about education. See Section 1.2.1 for more details.

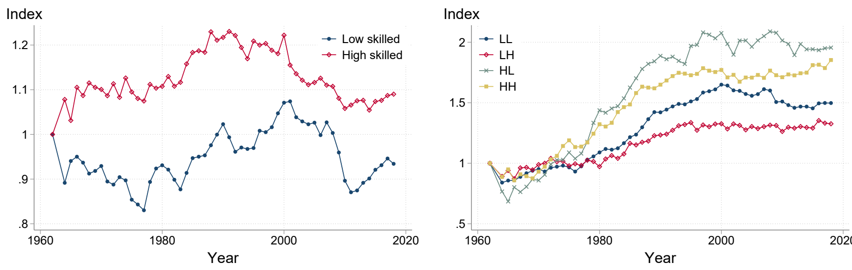

(a) *Single women, indexed* ***(b)*** *Women in couples, indexed*

Figure B.2: *Market hours among women in different types of households (indexed)*

Note: *Panel (a)*: indexed series of the lines in Figure 1.7a; *Panel (b)*: indexed series of the lines in Figure 1.7b.

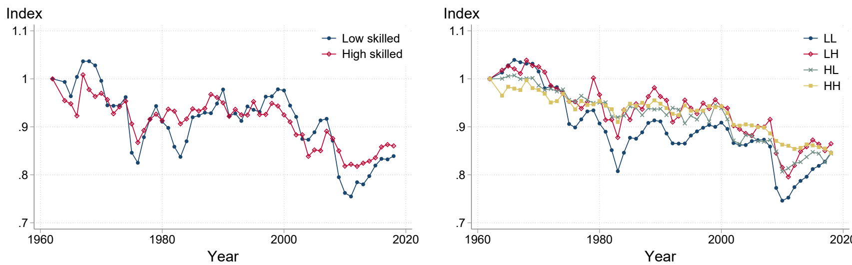

(a) *Single men, indexed* ***(b)*** *Men in couples, indexed*

Figure B.3: *Market hours among men in different types of households (indexed)*

Note: *Panel (a)*: indexed series of the lines in Figure 1.9a; *Panel (b)*: indexed series of the lines in Figure 1.9b.

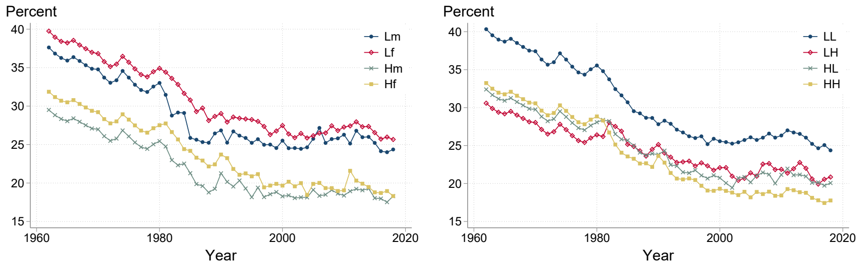

(a) Single households *(b) Couple households*

Figure B.4: *Expenditure share on nondurable goods*

Note: *Panel (a)*: Expenditure shares on nondurable goods among single men and women, depending on skill levels. *Panel (b)*: Expenditure shares on nondurable goods among couple households, depending on the levels of education of the two spouses.

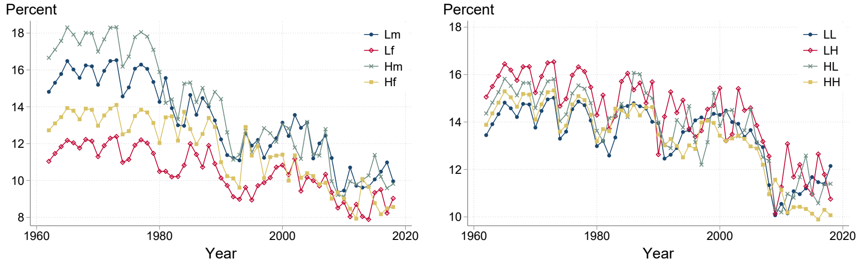

(a) Single households *(b) Couple households*

Figure B.5: *Expenditure share on durable goods/capital*

Note: *Panel (a)*: Expenditure shares on durable goods/capital among single men and women, depending on skill levels. *Panel (b)*: Expenditure shares on durable goods/capital among couple households, depending on the levels of education of the two spouses.

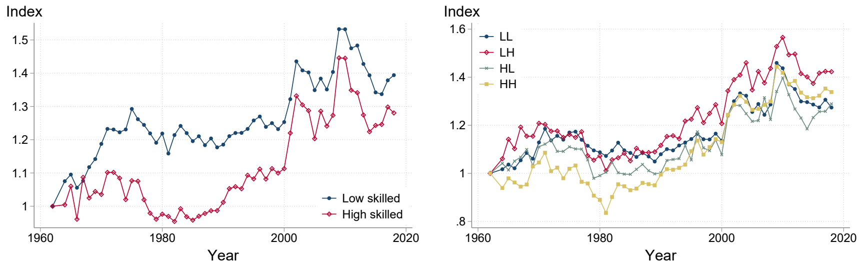

(a) Single women, indexed *(b) Women in couples, indexed*

Figure B.6: *Wage rates for women in different types of households (indexed)*

Note: *Panel (a)*: indexed series of the lines in panels Figure 1.14a; *Panel (b)*: indexed series of the lines in panels Figure 1.14b.

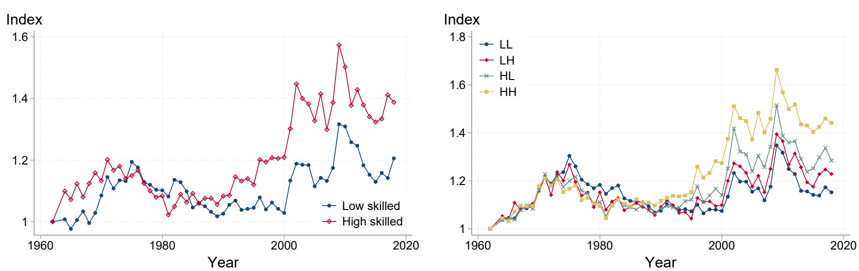

(a) Single men, indexed *(b) Men in couples, indexed*

Figure B.7: *Wage rates for men in different types of households*

Note: *Panel (a)*: indexed series of the lines in panels Figure 1.16a; *Panel (b)*: indexed series of the lines in panels Figure 1.16b.

Chapter 2

Home Production, Expenditures, and Time Use*

*I am grateful to Timo Boppart and Per Krusell for their advice and support. I would also like to take the opportunity to thank Mitch Downey, Richard Foltyn, Jose-Elias Gallegos, Stefan Hinkelmann, Philipp Hochmuth, John Kramer, Kieran Larkin, Kurt Mitman, Jonna Olsson, David Schoenholtzer, Josef Sigurdsson, Xueping Sun, and Has van Vlokloven.

2.1 Introduction

The shift over the last six decades in how households in the U.S. allocate expenditures was substantial. Figure 2.1 shows the ratios between quantities of services and nondurable goods, for low and high skilled single men and women, respectively, together with the relative price, between the years 1962 and 2018.[1] Despite the increase in the relative price, there is a clear and quantitatively large increase in the ratio up until the 1990s. The other pronounced pattern in the data is that the ratios for high skilled households are significantly higher at any given point in time.

Explaining these facts with traditional economic theory seems challenging. In particular, a standard framework with stable homothetic preferences implies: (i) that the ratio between quantities of services and nondurable goods falls when the relative price of services increases, and (ii) that the expenditure shares on different categories of consumption only depend on relative prices, and that, therefore, there cannot be any systematic differences between how different types of households allocate expenditures. Clearly, the variation across time, and across types of households, that is implied by the standard framework with stable homothetic preferences is contradicted by the data. A second salient fact about the data is the noticeable nonmonotonic pattern of the ratio. Any theory that seeks to explain why the ratio increased during the first three decades should also be able to account

[1] I focus on single men and single women households in this paper. From here on, I will refer to them as men and women, respectively. Details about how all data used in the current paper is collected and processed can be found in Almgren (2023). Data that is used in the current paper is smoothed using the HP-filter. I am more interested in low-frequency movements, such as changes over a decade, than in changes from year to year. Smoothing the data makes it easier to detect these types of trends.

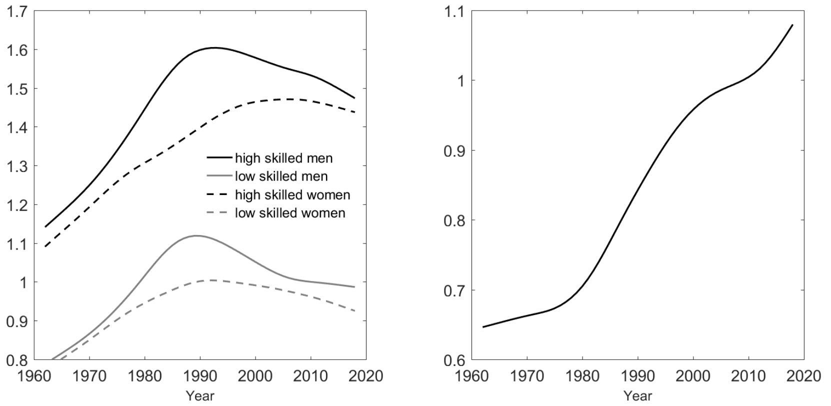

(a) *Serv.-to-nondurable goods ratio* ***(b)*** *Relative price of services*

Figure 2.1: *Services-to-nondurable goods ratio and the relative price*

Note: *Panel (a)*: Development of the services-to-nondurable goods ratio for high and low skilled men and women, respectively. The values are indexed to the value for low skilled men in the year 2010. *Panel (b)*: The relative price between market services and nondurable goods. The data is filtered using the HP-filter. For more details about the data, see the main text. Figure B.1 shows the unfiltered series for the relative price. The unfiltered series for the data in panel (a) are shown later, when compared to model output; see Sections 2.5 and 2.7.

for the subsequent decline.

Models with endogenous labor supply typically focus on hours worked in the market and view the residual as leisure. But there is no doubt that a significant fraction of hours that are not spent in the market are instead spent on home production (see, e.g., Ramey, 2007). If a substantial part of households' resources is used to produce output at home, it is necessary to consider home production when seeking to understand expenditures as well as market hours.

In this paper, I test whether a theory with stable homothetic preferences and home production has the potential to match the data. This includes specifying a set of functions and estimating parameter values. Households receive utility from consuming market and home produced services. The home production technology combines durable and nondurable goods, as well as hours in home production and outputs home services. Nondurable goods are often modeled in the literature as a final expenditure that is directly consumed (e.g., Bridgman, 2016; Moro et al., 2017; Olovsson, 2009, 2015). I choose to include it as an input in home production, because in the market, the same types of nondurable goods would often be bundled with labor and capital to produce market services. E.g., in the year 2018, 35 percent of nondurable expenditures were allocated to *Food and beverages purchased for off-premises consumption.* All items in this category, like beef, milk, or wine, just to name a few, would be combined with both labor services and capital by a restaurant or a cafe to produce a market service.[2] The home production technology uses standard functional forms (CES), but the particular way in which inputs are combined has not previously been used in the literature.

In the time-use dimension, choices are made about how many hours to supply to the market and how many hours to spend in production at home. Households face type-specific hourly wage rates, and goods and services in the market are bought at prices that are constrained to equal those in the data. The current paper hence makes no effort in trying to understand drivers on the supply-side that generate changes in the relative prices of market produced goods and services,

[2] I measure expenditures on market services, nondurable goods, and durable goods, respectively, using the final expenditures approach. An alternative is to consider how much of value added that is generated by each sector, i.e., the value added approach. For a discussion about these two alternatives, see Herrendorf et al. (2013).

or growth in wage rates over time.

The values that are estimated for the elasticities are important for the model predictions. One parameter governs the elasticity of substitution between market and home produced services, while two other parameters govern the elasticities of substitution in home production. I take into account that how much each unit of input in home production contributes to output depends on how efficiently it is used. For each type of input, the efficiency can vary across types of households and change over time. My method for estimating the input-specific productivity growth rates is similar to that in Katz and Murphy (1992). In their case, they need to assume, similar to what I do, a production function and a value for the elasticity of substitution between high school and college workers to estimate what they refer to as ”relative demand shifts”, which potentially captures skill-biased technical change. A similar strategy is also used by Hassler et al. (2021).

The framework is successful in that it can account for most of the variation over time and across types of households, in terms of how they allocate time and expenditures. The list of outcome variables is not limited to the ratio between the quantities of market services and nondurable goods but, as a standard framework with stable homothetic preferences cannot account for how it varies across time and space, I put extra emphasis on it. I now summarize some findings on time-use and the ratio between market services and nondurable goods, respectively, starting with the latter. The presentation and analysis of other results is deferred to later.

To understand the development of the ratio between the quantity of services and nondurable goods, it is useful to decompose it into two components. The first component captures substitution between market and home produced services, and the second captures substitution between different inputs in home production.

The cost of producing at home is captured by the *implicit price of home production*, which I derive.[3] An ex-ante plausible story is that the increase in the services-to-nondurable goods ratio, that was observed up until the 1990s, was caused by an increase in the implicit price of home services, perhaps driven by an increase in the wage rate, and that households substituted into market services. With respect to men, my results tell something different. I find that the implicit price of home production fell sharply up until the 1990s, and from that point in time it decreased only little. Given homothetic preferences, this unambiguously led men to substitute away from market services into home services – which had a negative effect on the ratio in question.

The fall in the implicit price of home production was caused by falling prices per efficiency units of durable and nondurable goods, respectively. These, in turn, were partly driven by decreases in the prices of durable and nondurable goods, respectively, relative to the price of market services. In part they were also driven by positive changes in the productivities that augment each input into its efficiency unit counterpart.

It was the fact that men substituted away from nondurable goods in home production that led to the increase in the market services-to-nondurable goods ratio. There are in particular two parameters that together drove substitution in this direction. First, for the first

[3]For men, this cost is measured in terms of foregone units of market services. The implicit price is constructed slightly differently for women, however, as it accounts for social norms. How the implicit price is defined does not affect the series for expenditures and time that the model produces. It does, however, aid in analysing the changes over time.

three decades I find sizeable growth in the efficiency term that augments nondurable goods in home production. Second, I find the respective inputs in home production to be gross complements. With the inputs being gross complements, increases in the efficiency term on nondurable goods made men substitute to other inputs in home production, which, hence, was the major driver behind the increase in the ratio between the quantity of services and nondurable goods.

From the beginning of the 1990s, the implicit price of home services decreased marginally, leading to some substitution away from market services to home services. At the same time, men increased the quantity of nondurable goods used in home production, driven by a decline in the relative price of nondurable goods, but also by the fact that the nondurable goods augmenting technology stopped increasing. The latter channel (substitution in home production) dominated the first (substitution between market produced and home produced services), leading the market services-to-nondurable goods ratio to fall.

For women, the story is somewhat different. Through the lens of my theory, the behavior of women can be understood as being guided in part by a time-varying effort cost of working in the market relative to the time-varying effort cost of working in home production. I refer to this as changes in "social norms", capturing the natural phenomenon that women may have felt less and less discriminated at work over time, perhaps complemented with the fact that the gradual growth of the service sector offered more and more attractive occupations for women. In contrast to the development for men, changes in social norms increased the implicit prices of home services up to around the year 2009 and made women substitute away from home services to market services. After the year 2009, social norms did only have minor effects on the development, while other factors caused a small decrease in the ratio between market services and home services. A decrease in the importance of nondurable goods in home production also contributed positively to the increase in the ratio between market services and nondurable goods up until the early 1990s. Like for men, more nondurable goods were used in home production after that point in time, primarily driven by the fact that the nondurable goods augmenting technology stopped growing.

With respect to the time-use for low and high skilled men, the model can match the fact that market hours decreased, hours in home production increased, and total hours worked (i.e., the sum of hours in home production and market hours) decreased over time. The decrease in total hours worked was generated by the fact that the implicit real wage[4] increased over time together with the fact that the income effect dominates. Considering that total hours fell, the increase in the hours in home production came from households re-allocating hours from the market to home production. This was mainly driven by home services becoming cheaper to produce relative to the price of market services.

Among women, the time allocation did especially change among the low skilled. Total hours worked decreased by around eight hours between the years 1962 and 2018. The model does, however, falsely generate that about half of this decline came from decreasing market hours, while in fact is was only about one hour. Qualitatively, the direction of the changes generated by the model for market hours

[4]The implicit real wage is defined as the real wage (which in turn is the nominal wage rate normalized with respect to the price of market services) divided through by the implicit price index of the consumption index C, which is a composite with market and home services. See Section 2.3 where the model is presented, but also Section 2.5.2 for more details about C and the implicit price.

and hours in home production aligns with developments in the data, and the model explains almost half of the total decline in hours in home production between the years 1962 and 2018. A key result in this paper is that hours worked at home would not have decreased over time, but would instead have increased, had social norms not changed. Absent changes in social norms, market hours would have decreased.

The rest of the paper is organized as follows. Section 2.2 discusses some literature connected to the current paper. In Section 2.3, I present the model, and then, in Section 2.4, I go through how parameters are estimated using data for men. The results for men are presented in Section 2.5. Sections 2.6 and 2.7 proceed with estimation and results for women. The importance of social norms is explicitly considered in Section 2.8. Section 2.9 then concludes the paper.

2.2 Connections to the Literature

First, this paper relates to the broad field of economics that researches the drivers behind structural change, more specifically the reallocation of economic activity across agriculture, manufacturing, and services (Herrendorf et al., 2014). Typically, reallocation originates from changes in relative prices, increases in real income, or both. As production becomes more efficient over time, real income rises. Changes in relative prices can be generated in different ways. Acemoglu and Guerrieri (2008) generate relative price changes and reallocation across sectors by assuming different capital intensities, while Caselli and Coleman II (2001) argue that a relative increase of skilled workers decreased the relative prices of non-agricultural goods. A more common assumption is that it is differential productivity growth rates across the sectors that cause changes in relative prices, as in, e.g., Ngai and Pissarides (2007).

While increases in income cause sectoral reallocation via income effects when preferences are nonhomothetic, the latter cause reallocation via the substitution effect, which is the case also under the assumption that preferences are homothetic. Kongsamut et al. (2001) assume identical growth rates in all sectors and hence all model-generated reallocation between sectors is driven by income effects. Ngai and Pissarides (2007), on the other hand, allow for differential growth rates but use homothetic preferences, thereby shutting down income effects. While each of these can account for some stylized facts about structural change, neither of them can match them all.

With those results in mind, many papers allow for both income and substitution effects, e.g., Boppart (2014) specifies an indirect utility function with these properties. He successfully accounts for, e.g., the fact that the service-to-goods ratio increased during many decades despite an increase in the relative price between services and goods, and why higher-income households systematically allocate a larger fraction of expenditures to services. Building on the same type of indirect utility function, Cravino et al. (2022) find that about a fifth of the increase in the service expenditure share in the U.S. between 1982 and 2016 can be explained by population aging. Comin et al. (2021) introduce nonhomothetic constant elasticity of substitution preferences into a model that otherwise closely follows workhorse models of structural transformation (e.g., Herrendorf et al., 2013). They find that three fourths of sectoral reallocation can be accounted for by income effects, while the remaining fourth can be accounted for by substitution caused by changes in relative prices. This echoes the finding of Herrendorf et al. (2013), who also find income effects to be most important. Hubmer (2023) investigates the stability in the labor share in the U.S. up until around the early 1980s and why it fell thereafter. He estimates capital and labor to be gross substitutes in production and that technology growth was capital biased before the early 1980s. However, this effect was balanced by an overall increase in real income, which caused households, via non-homothetic preferences and the income effect, to demand more labor intensive goods (he refers to it as goods, but "goods" that are produced in labor-intensive sectors often tend to be services). After the early 1980s, technology growth became more capital-biased and this caused the labor share to decrease.

Ever since Becker (1965) urged economists to focus more on how time is allocated outside work, home production has become more important in many strands of the economics literature. Starting decades back, home production has helped many models to successfully explain various facts about how, e.g., output, investment, and market hours vary (and co-vary), see, e.g., Benhabib et al. (1991); Greenwood and Hercowitz (1991); Greenwood et al. (1995); McGrattan et al. (1997); Einarsson and Marquis (1997); Baxter and Jermann (1999); Gomme et al. (2001) and Gomme et al. (2004). Empirical papers also find substitution between market hours and hours in home production to be important, see, e.g., Aguiar et al. (2013) and Burda and Hamermesh (2010), who document that a significant fraction of variations in market hours over the business cycle are absorbed by home production, and Leukhina and Yu (2021) who find that much of the decline in market hours during the COVID-19 recession was absorbed by home production.

Home production has also been introduced to highlight various channels that have earlier been overlooked. Just to name a few, Boerma and Karabarbounis (2021) find that welfare inequality is amplified in their framework with home production, and Calvo et al. (2021) argue that between-household inequality increases as a result of the spouse's hours being complements in home production, which generates positive marriage sorting. Home production is used by Aruoba et al. (2016) to explain the empirically observed relationship between housing wealth and nominal interest rates.

More closely related is the literature that uses home production to understand reallocation across sectors in terms of production and consumption. Rogerson (2008) and Ngai and Pissarides (2008) are examples that explain shifts in labor between different sectors in the economy by assuming that productivity grows at different rates between different sectors, and between home and the market. These papers focus on production in the market economy, while I focus more on home production and the demand side. Moreover, Rogerson (2008) makes use of non-homothetic preferences while mine are homothetic.

Home production has been used to understand differences in time-use across countries, and developments over time. Olovsson (2009) compares hours in home production and in the market in the U.S. and Sweden, respectively. Using a model that includes home production, he finds that almost all of the differences can be accounted for by differences in taxes.[5] Duernecker and Herrendorf (2018) present somewhat contradictory results that point to that increases in labor taxes lead households to increase leisure, but that hours in home production are unaffected. They show that this can happen in a model with home production and nonhomothetic preferences when the growth in

[5] He also derives, in Olovsson (2015), optimal tax rates for the U.S. and Sweden, using a similar framework that includes home production. Since home production cannot be taxed directly, he finds that the tax rates on market services are too high.

home labor productivity is strong. Work by Greenwood et al. (2005) and Greenwood et al. (2016) suggests that a large part of the substantial increase in the labor force participation rate seen in the U.S. was driven by a sharp decrease in the price of capital. As capital is assumed to be a factor of production at home together with labor, and with labor and capital being gross substitutes in their cases, this freed up time for women that they could instead use to do market work. Although my setting is vastly different, the fall in capital prices will indeed play a role in my framework. In Fang and Zhu (2017) it is shown that a model with home production can explain a big part of the cross-sectional variation and developments in time-use over time. They focus on married households in which both spouses do market work, while I focus on single households and also consider the extensive margin.[6] Jones et al. (2015) argue through the lens of their model, which considers home production, that the decrease in the gender wage gap was important for the increase in the labor force participation rates among married women, but had essentially no effects on the participation rates among single women.

Boppart and Ngai (2021) use home production to understand trends in leisure. Growth in wage rates up until the early 1990s increased leisure for all skill types, as consumption and leisure are gross complements (channel one). After the early 1990s, the trends in leisure diverged and for high skilled households leisure has been trending down. As the wage rates were expected to grow faster for highly educated individuals, this meant that the implicit price of leisure was expected to become more expensive in the future. Households therefore chose to instead increase leisure today (channel two). Channel

[6] I do not model or target the decision about whether to work in the market or not. However, the time-use data used in the current paper is based on averages that account for extensive margin choices. See Almgren (2023) for more details.

two was dominating for highly educated individuals after the early 1990s, while leisure time for individuals with lower levels of education continued to increase because the first channel still dominated. It has also been argued by Rupert et al. (2000) that not taking home production into account might lead to severely underestimating the intertemporal substitution elasticity.

A closely related paper by Moro et al. (2017) finds that home production is by itself not sufficient to be able to account for shifts between sectors in the economy, and needs to complement it with non-homothetic preferences. My framework is different from theirs in many respects and I will mention some of these. In their setup, nondurable goods are consumed directly and give instant utility. My framework instead models nondurable goods as an input in home production, where it is combined with a composite of durable goods and hours in home production to produce output. Second, I consider factor-specific technological change, while they introduce factor-neutral technical change in home production. Third, total hours worked in the market and at home are taken as given in their model, while in my framework the household chooses the total amount of time that is to be dedicated to work, as well as how the time should be distributed between the two different work activities. Fourth, while they consider one representative household, I look at different types of households, grouped by gender and skill level.

Non-homothetic preferences in Buera and Kaboski (2012a) results in households demanding more services as productivity and income increase. Services can either be produced in the market or at home. Individuals can specialize in their production, but this will only affect their productivity in the market. As more demand is directed to services, more individuals will choose to specialize and work in the market to satisfy said demand, which will cause the market service sector to grow. In a companion paper, Buera and Kaboski (2012b) employ the same utility function but instead focus on new technology in the form of capital having the potential to make production significantly cheaper. New types of technologies are generally expensive. They argue that firms usually are the ones that first purchase and employ the new technology, as their scale makes it profitable. As a result of firms applying the technology, their production costs and prices towards households fall, which leads to marketization. Over time, the price of capital falls drastically and it becomes affordable for households to purchase and use it in home production. This potentially leads home services to first be marketized but then de-marketized later on. They document that historically, many types of services have, in agreement with their prediction, followed this cycle. Interestingly, my results go in a similar direction, as I find that men substitute away from market services to home produced services. Although this is not only driven by a decrease in the price of capital, it is definitely important.

2.3 Model

In this section, I lay out a theoretical model. I focus on households' decisions on time-use and how to allocate expenditures. Households do not enter with any assets. The intertemporal consumption-savings choice is not modeled (see, e.g., Duernecker and Herrendorf, 2018; Hubmer, 2023, who also repeatedly solve static models). Wage rates, and prices of goods and market services, respectively, are fed in exogenously and are taken as given by households.

Utility is derived from consuming market and home produced services. Market services are, as suggested by their name, purchased in the market, and consumed directly by households. Home services, on the other hand, cannot be traded in the market and hence need to be both produced and consumed by the household itself. To produce these home services, the household combines its own hours with durable and nondurable goods in a nested CES production function. The particular way in which I model home production has not previously been considered in the literature. Individuals supply labor to the market and work at home. Both of these activities generate disutility. I focus on single households, which are divided into four groups, according to their gender (man/woman) and skill level (high/low)[7]. The four different types of households are hence the four possible combinations of gender and skill level.

Research has found social norms to be important to explain historical developments in the female labor participation rate (see, e.g., Lee, 2021; Eckstein and Lifshitz, 2011; Fernández, 2013).[8] See also Balleer et al. (2014), who find that the age-participation profiles for women in six European shifted up between 1983 and 2007, and argue that social norms are likely to explain that development. My model includes parameters that are intended to capture the development of such norms. More specifically, the parameters affect women's disutility from market work and work in home production and are allowed to change over time.

Like when the market produces something, I consider that there can be technological change in home production as well. Technological change is factor specific; there is one parameter for each of the

[7] I do as in Almgren (2023) and classify an individual as high skilled if he/she has at least a college degree, and otherwise as low skilled.

[8] The literature most often focuses on explaining the increase in labor force participation rates among married women. Although I do not focus on married women here, I find it likely that norms about women's participation in the labor market in general, i.e., also for single women, have changed over time.

three different production factors that are combined to produce home services.

2.3.1 The Household Problem

The four types of single households are indexed by $j = \{LM, LF, HM, HF\}$, where the first letter indicates skill level: *L* for low and *H* for high, and the second letter indicates the sex: *M* for man and *F* for woman.

A household faces four explicit prices: the price of market services (which serves as the numeraire in each period), the price of nondurable goods, P_n, the price of durable goods/capital, P_k, and an hourly wage rate, W. While the hourly wage rate is household type specific, the other prices are the same for all types. The household chooses how many hours to work in the market, L_m, and hence earns income WL_m.

Expenditures are allocated between market services, C_m, and two kinds of goods: nondurable, N, and durable, K. The two types of goods are not consumed immediately, but are used as factors of production, or as inputs, in home production. To produce an output at home, which I will refer to as home services, C_h, the household also needs to supply its own labor, L_h. The home production technology is a nested CES: the inner nest combines hours in home production with durable goods into an index Q and this index, in turn, is combined in an outer CES with nondurable goods to produce home services. The factor augmenting technologies are z_l, z_k and z_n.

The household has preferences over both market and home services and these are also combined in a CES function. There is no restriction on how many hours that the household can allocate to market and home production, respectively, but both of the activities are associated with disutility. Parameters κ_m and κ_h represent social norms and for men these are equal to one by assumption. The model will be used to generate household type specific time series related to consumption and time allocation. I do not model the intertemporal consumption-savings choice, but instead solve the static model repeatedly, for each year. Subscripts j and t are used to indicate which prices and parameters that can vary across households types and time, respectively. In each period t, a household of type j solves a constrained maximization problem, which can be formulated as follows:

$$\begin{aligned} \max \quad U_{jt} = \frac{1}{1-\gamma}\left(C_{mjt}^{\frac{\sigma_{mh}-1}{\sigma_{mh}}} + C_{hjt}^{\frac{\sigma_{mh}-1}{\sigma_{mh}}}\right)^{\frac{\sigma_{mh}(1-\gamma)}{\sigma_{mh}-1}} \\ - \chi_j \frac{\left(\kappa_{mjt}L_m + \kappa_{hjt}L_h\right)^{1+\varphi}}{1+\varphi} \end{aligned} \tag{2.1}$$

$$\text{where} \quad C_{hjt} = \left((z_{njt}N_{jt})^{\frac{\sigma_{nq}-1}{\sigma_{nq}}} + Q_{jt}^{\frac{\sigma_{nq}-1}{\sigma_{nq}}}\right)^{\frac{\sigma_{nq}}{\sigma_{nq}-1}} \tag{2.2}$$

$$\text{and} \quad Q_{jt} = \left((z_{kjt}K_{jt})^{\frac{\sigma_{kl}-1}{\sigma_{kl}}} + (z_{ljt}L_{hjt})^{\frac{\sigma_{kl}-1}{\sigma_{kl}}}\right)^{\frac{\sigma_{kl}}{\sigma_{kl}-1}} \tag{2.3}$$

$$\text{subject to} \quad W_{jt}L_{mjt} = P_{nt}N_{jt} + P_{kt}K_{jt} + C_{mjt} \tag{2.4}$$

It seems plausible that an individual who works long hours in the market has a greater distaste for spending hours in home production when returning from market work, than an individual who worked few hours in the market. The specification results in that the marginal disutility of hours worked in home production depends also on the amount of hours supplied in the market (and vice versa). Note lastly that $\kappa_{mjt} = \kappa_{mt}$ and $\kappa_{hjt} = \kappa_{ht}$ $\forall j \in \{LF, HF\}$.

I use the model to produce time series, by feeding in sequences of wages and prices $\{W_{jt}\}_{t=1962}^{t=2018}$, $\{P_{nt}\}_{t=1962}^{t=2018}$ and $\{P_{kt}\}_{t=1962}^{t=2018}$. The absence of an intertemporal margin motivates the question why I need both parameters γ and φ, especially since these can not be separately pinned down. I choose to include γ as a free parameter, in order to be able to match trends in time-use with a positive value for φ. Without γ, the parameter φ would have to capture the convexities in the disutility of working as well as any concavities in the utility of consumption. That said, it would be possible to match the trends in hours even without γ, *if* I allowed $\varphi < 0$.

2.4 Estimation: Men

Given the parameter values, the model makes quantitative predictions, that can be categorized into two groups. The first is about time-use. Households decide how many hours of leisure to have, and then how to allocate working hours between the market and their home. This is in contrast to most papers in the literature with home production, which often restrict a household's time allocation decision, e.g., by fixing the amount of leisure. Given the number of hours that the household spends working in the market, it will receive an income. All this income is then spent, and the household decides how to allocate it between capital, nondurable goods and market services. All these decisions depend on wage rates and relative prices.

Predictions about time-use and allocations of expenditures will be compared to data realizations, for the different types of households and across time. The underlying time series are constructed in the companion paper Almgren (2023), so for any details about the data, I refer the reader there.

Although the model parameters are many, they are greatly outnumbered by the number of data moments. One part of the strategy for how the parameter values are estimated, is to start with men, and then move on to women. Each step pins down the values of some of the model parameters, which are then taken as given in the next step.

If the part of the utility function that measures disutility from hours worked is disregarded, it can be seen as containing only a nested CES with three layers. The innermost layer combines capital and hours in home production into an index Q and this index, in turn, is combined in the mid layer with nondurable goods to produce home services. In the outer layer, home services are combined with market services into an index that can simply be denoted by C. The procedure for estimating the parameters involves starting with the equations that relate directly to the innermost nest. Thereafter, I progress outwards.

The process can be divided into two stages: the first stage relates directly to how expenditures are allocated and how the household substitutes between different inputs in home production. I search over combinations for the parameters that govern the elasticities of substitution: σ_{kl}, σ_{nq}, and σ_{mh}. Given each combination, I identify the values for the other parameters that are needed to match the data in the cross-section as well as over time. In the second stage, I focus on parameters that pin down the levels of hours and expenditures.

2.4.1 Substitution Between Home Hours and Capital

I start with the household's optimal choice about how many units of L_h relative K to use. This is given by the following equation:

$$\frac{L_{hjt}}{K_{jt}} = \left(\frac{P_{kt}}{W_{jt}}\right)^{\sigma_{kl}} \left(\frac{z_{kjt}}{z_{ljt}}\right)^{1-\sigma_{kl}} \tag{2.5}$$

z_k and z_l are capital and hours in home production augmenting technology parameters, respectively. The ratio between hours and capital increases in the relative price of capital and is also affected by the ratio between capital and home labor augmenting technology. Positive changes in z_k/z_l make the household increase hours in home production relative to capital if hours and capital are gross complements in home production.

At this point, I consider values of σ_{kl} and for each value I solve for the path of z_k/z_l that brings the model closest to the hours in home production-to-capital ratio in the data. Here, it is only the ratio z_k/z_l that can be pinned down (i.e., for each value of σ_{kl}). I allow for the ratio to vary across household types. However, if expressed as an index, the paths for z_k/z_l are restricted to be the same for all household types. The path for z_k/z_l that is attached to each elasticity parameter minimizes the sum of mean squared errors for the two types of men.

2.4.2 Substitution Between Nondurable Goods and Q

The middle nest combines the index Q with nondurable goods, N, to produce home services, C_h. The optimal ratio for nondurable goods relative to the composite Q is given by

$$\frac{N_{jt}}{Q_{jt}} = z_{njt}^{\sigma_{nq}-1}\left(\Omega_{jt}^{\frac{1}{1-\sigma_{kl}}}\frac{P_{kt}}{z_{kjt}}\right)^{\sigma_{nq}} P_{nt}^{-\sigma_{nq}}, \text{ where } \Omega_{jt} = 1 + \left(\frac{z_{ljt}P_{kt}}{z_{kjt}W_{jt}}\right)^{\sigma_{kl}-1} \tag{2.6}$$

The composite Q is not directly observed, however. Q is constructed by K and L_h (see Equation 2.3), and can, using Equation (2.5), be written as

$$Q_{jt} = \Omega_{jt}^{\frac{\sigma_{kl}}{\sigma_{kl}-1}} z_{kjt} K_{jt} \tag{2.7}$$

Combining Equation (2.6) with Equation (2.7) yields

$$\frac{K_{jt}}{N_{jt}} = \Omega_{jt}^{\frac{\sigma_{nq}-\sigma_{kl}}{\sigma_{kl}-1}} \left(\frac{z_{kjt}}{z_{njt}}\right)^{\sigma_{nq}-1} \left(\frac{P_{nt}}{P_{kt}}\right)^{\sigma_{nq}} \tag{2.8}$$

At this stage, I consider combinations for $(\sigma_{kl}, \ \sigma_{nq})$. Each value for σ_{kl} comes with paths for the ratio z_k/z_l, as described in Section 2.4.1. From here, the procedure is conceptually the same as in the previous section: a path and levels for z_k/z_n are attached to each pair $(\sigma_{kl}, \sigma_{nq})$. Like earlier, the relative changes over time in z_k/z_n are restricted to be the same for all households, but the levels of the ratios are allowed to be different.

2.4.3 Substitution Between Market Services and Home Services

The ratio between home services and market produced services is given by

$$\frac{C_{hjt}}{C_{mjt}} = \Gamma_{jt}^{\frac{\sigma_{mh}}{\sigma_{nq}-1}} \left(\frac{z_{njt}}{P_{nt}}\right)^{\sigma_{mh}}, \text{ where } \Gamma_{jt} = 1 + \Omega_{jt}^{\frac{\sigma_{nq}-1}{\sigma_{kl}-1}} \left(\frac{z_{kjt}P_{nt}}{z_{njt}P_{kt}}\right)^{\sigma_{nq}-1} \tag{2.9}$$

As C_h is not observable, it is expressed as a function of N. Combining Equation (2.2) with equations (2.7) and (2.8) yields

$$C_{hjt} = \Gamma_{jt}^{\frac{\sigma_{nq}}{\sigma_{nq}-1}} z_{njt} N_{jt} \tag{2.10}$$

Combining equations (2.9) and (2.10) then gives

$$\frac{C_{mjt}}{N_{jt}} = \Gamma_{jt}^{\frac{\sigma_{mh}-\sigma_{nq}}{1-\sigma_{nq}}} P_{nt}^{\sigma_{mh}} z_{njt}^{1-\sigma_{mh}} \tag{2.11}$$

At this stage, I evaluate combinations of $(\sigma_{kl}, \sigma_{nq}, \sigma_{mh})$. Remember that combinations of σ_{kl} and σ_{nq} already are already attached with values on factor augmenting productivity parameters, such as paths and levels for z_k/z_n. Γ is independent of σ_{mh} and hence predetermined. For each combination $(\sigma_{kl}, \sigma_{nq}, \sigma_{mh})$, this stage solves for household-specific levels and common (indexed) paths of z_n, such as to minimize the mean squared errors. Once z_n has been determined for each household, this also pins down the levels of z_k and z_l. When I later on focus on women, a combination $(\sigma_{kl}, \sigma_{nq}, \sigma_{mh})$ is attached with how z_l, z_k and z_n, respectively, grow over time.

2.4.4 Parameters That Govern Levels, and Pinning Down Parameter Values

Each combination of elasticities $(\sigma_{kl}, \sigma_{nq}, \sigma_{mh})$ is coupled with factor augmenting technologies (z_l, z_k, z_n) that can change over time and differ between the two types of men. These parameters together determine the model-implied consumption ratios and expenditure shares, but the model is still silent about levels. Before going into the next step, I discard some combinations that have been produced up to this point.[9]

[9] Combinations are discarded if the growth rate for any of the factor augmenting technologies is considered too low or too high, and/or the implied cross-sectional differences in at least one of the respective technologies is considered too high. In particular, I discard a combination if the technology either falls between the first and final period, or is higher in the last period than what would be the case with a cumulative average growth rate of four percent. The restriction put on the

The next step involves pinning down the values for elasticities and factor augmenting technologies, which is done while also pinning down values for the remaining three parameters γ, φ and χ. These influence the levels of hours worked and expenditures. In order to match cross-sectional level differences in time-use, χ can vary between households.

Let σ be a $G \times 3$ matrix with all the remaining G combinations of elasticities that remain. Each combination $g \in G$ comes with given indexed paths and levels for the factor augmenting technologies. For each combination g, I solve for the values of φ, γ, χ_L, and χ_H, that together minimize the sum across household types of the mean squared errors for time-use. Specifically, the targeted data points are hours worked in home production and work in the market, respectively, in each year for both low skilled and high skilled men.

I then choose the combination g^* that satisfies the following:

$$B_{g^*} \leqslant B_g \; \forall g \in G, \text{ where } B_g = E_g + T_g + (1 - R2_g) + Z_g \tag{2.12}$$

where E, T, $R2$, and Z are all $G \times 1$ vectors. The first vector E contains model errors (differences between model predictions and data) related to the ratios L_h/K, K/N and C_m/N. Specifically, it is the mean squared errors of each of these respective ratios in each year, averaged across the two household types. The next vector, T, summarizes the model errors for time-use in a similar way, by for each type computing the mean squared errors for hours in home production and market hours and then averaging across the two household types. I add a third criterion that induces a preference for more linear paths for the factor augmenting technologies. This statistic is computed by, using OLS, fitting a linear line to each of the indexed paths and then

cross-sectional differences is that for neither of the factor augmenting technologies, one household type can be more than five times as productive as the other.

computing the R^2. For each combination g I get three measures of R^2, one for each specific technology, denoted by R_l^2, R_k^2, and R_n^2. These are mapped into $R2$ by summing the squares of each respective R^2: $R2 = (R_l^2)^2 + (R_k^2)^2 + (R_n^2)^2$. Since higher values for $R2$ are preferred, $(1 - R2)$ enters the objective function. Less cross-sectional dispersion in the factor augmenting technologies is preferred and this is achieved by having Z in the objective function:

$$Z_g = \left(\frac{z_{gna}}{z_{gnb}}\right)^2 + \left(\frac{z_{gka}}{z_{gkb}}\right)^2 + \left(\frac{z_{gla}}{z_{glb}}\right)^2 \tag{2.13}$$

A few comments about this is in order. A smaller value of Z_g is preferred over a larger value. For the entries on the right-hand-side of Equation (2.13), the subscripts on the zs include a or b. This is simply because it matters which of the two households that is placed on top. Take z_n, for example. If it is the case that $z_{n,high} >= z_{n,low}$, then $z_{na} = z_{n,high}$ and $z_{nb} = z_{n,low}$. Otherwise, low and high skilled switch places. This logic is applied to all the other ratios as well. The ratios are raised to the power of 2, which limits how the objective function will trade off an increase in one ratio by an increase in another. There is no need to consider the ratios in any particular year, since the cross-sectional variation is restricted to be constant over time.

While E is based on model-vs-data differences of ratios, T is based on differences of levels. $R2$ and Z are statistics that are based on parameter values. The vectors that enter the objective function in Equation (2.12) are normalized by, for each objective, subtracting the mean and then dividing by the standard deviation. As an example, let $\tilde{E}$ be the vector with model-errors for the ratios, let $\bar{E}$ be the mean of this vector and sd_E be the standard deviation. Then $E = (\tilde{E} - \bar{E})/sd_E$

is the vector that enters the objective function. I put equal weights on all arguments.

2.5 Results: Men

2.5.1 Parameters

Elasticities of substitution

Estimated values for the elasticities are reported in Table 2.1. The elasticity of substitution between capital and hours in home production is 0.66. These factors of production are hence gross complements, which has implications for the story behind the developments on which this paper focuses. I will go into these details later. The value is lower than what Bridgman (2016) uses as a baseline value (1.37). In both Moro et al. (2017) and Olovsson (2015), the elasticity of substitution between hours and capital in the home production function is assumed to be one.

Table 2.1: *Estimated elasticites*

σ_{kl}	σ_{nq}	σ_{mh}
0.66	0.09	0.45

Labor and capital create the composite Q, which is then combined with nondurable goods to produce home services. I estimate that the households use Q and N as complements in home production. More precisely, the elasticity is estimated to be 0.09. Since the particular way in which I model home production has not been considered before in the literature, there is no direct way to compare this result.

The elasticity of substitution between home services and market services, σ_{mh}, is estimated to be 0.45. This value is significantly lower than what is used in Ngai and Petrongolo (2017) and Boppart and Ngai (2021) who use values for the elasticity of substitution between market and home produced services of 2 and 1.5, respectively. It is also lower than the value in Moro et al. (2017), who set it to 2.3. Olovsson (2015) also uses a larger value, at 2.5. There are many possible reasons for why it is different from what is usually found or used in the literature. One possibility is that while other papers typically focus on aggregates, which sum over men and women from various types of households, I estimate it using single men households. Another thing that is important to remember is that how home services are produced, and which factors of production and inputs that are used, is different from the literature, thus making comparisons difficult.

Factor augmenting technologies

Two things are of interest: (i) changes over time for each of the factor augmenting technologies, and (ii) cross-sectional differences across the two types of men. Results regarding the latter are shown in Table 2.2. High skilled men are estimated to be almost twice as productive at using nondurable goods in home production, while they are slightly less efficient than low skilled men at using capital. The largest cross-sectional difference is for hours in home production: high skilled men are about 2.6 times as efficient at using their hours as are low skilled men.

Table 2.2: *Estimated cross-sectional differences in factor efficiency*

$\frac{z_{n,high}}{z_{n,low}}$	$\frac{z_{k,high}}{z_{k,low}}$	$\frac{z_{l,high}}{z_{l,low}}$
1.95	0.85	2.61

The estimated indexed paths are shown in Figure 2.2. In total over the whole period, capital productivity has grown the most, and (home) labor productivity the least. Capital productivity grew steadily until the beginning of the 1990s, then decreased until the beginning of the 2000s, and has increased since then. The cumulative average growth rate for z_k has been 2%. Labor productivity fell slightly until the end of the 1970s and has increased thereafter. The efficiency at which households use nondurable goods increased almost as much as capital efficiency until the mid 1990s but has since then decreased marginally. The cumulative average growth rates for z_l and z_n have been 0.5% and 1%, respectively.

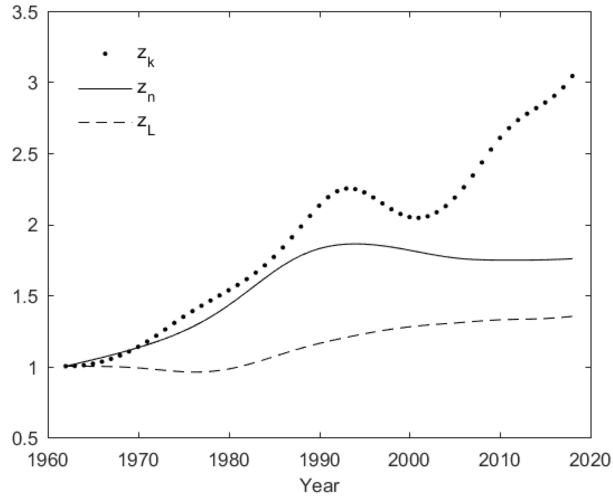

Figure 2.2: *Factor augmenting technologies*

Note: This graph illustrates the development of each factor augmenting technology between the years 1962 and 2018. Each series is indexed to its starting value in the year 1962. See the text for more details.

γ, φ, and χ

The value for γ is set to 1.3 and the estimated values of the remaining parameters are shown in Table 2.3. φ is estimated to be 1.82 and with the ratio $\chi_{\mathtt{high}}/\chi_{\mathtt{low}}$ being equal to 0.47, low skilled are estimated to have a larger distaste for work than do high skilled.

Table 2.3: *Estimated parameters related to labor supply*

φ	$\frac{\chi_{\mathtt{high}}}{\chi_{\mathtt{low}}}$
1.82	0.47

2.5.2 Model Generated Time Series

This section compares time series that are generated by the model with their respective data counterparts. I primarily focus on the moments that were used for the estimation but also consider other time series of interest. As the related literature often focuses on expenditure shares, I include figures that present how these develop in Appendix 2.B. I first focus on how expenditures are allocated and then move on to compare model-generated time series with data for time-use. There will be an ongoing discussion about the mechanisms at work. Only men are considered in this section. To save on notation, I only differentiate between household types by indicating their skill levels.

Ratios and expenditure shares

Hours in home production-to-capital ratio: The innermost nest in the home production function combines hours in home production with capital and this will be my first focus. The model-generated time series for men are compared with their data counterparts in Figure

2.3a. How does the model explain the data? In any given year, the cross sectional dispersion in the ratio is given by

$$\frac{L_{h,low,t}/K_{low,t}}{L_{h,high,t}/K_{high,t}} = \left(\frac{W_{high,t}}{W_{low,t}}\right)^{\sigma_{kl}} \left(\frac{z_{k,low,t}/z_{l,low,t}}{z_{k,high,t}/z_{l,high,t}}\right)^{1-\sigma_{kl}} \tag{2.14}$$

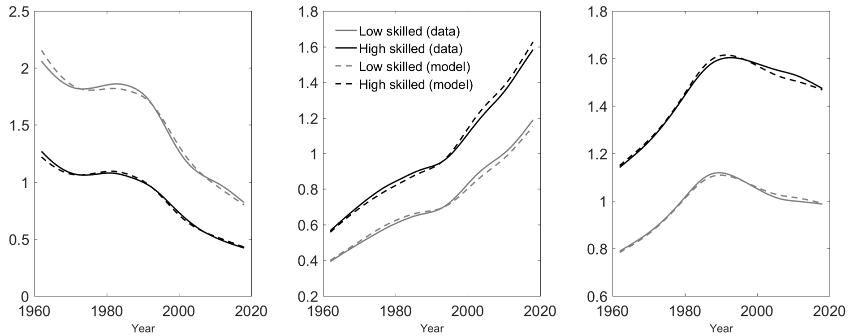

(a) L_h/K ***(b)*** K/N ***(c)*** C_m/N

Figure 2.3: *Quantity ratios for men: model and data*

Note: *Panel (a)*: Ratio between home hours and capital. *Panel (b)*: Ratio between capital and nondurable goods. *Panel (c)*: Ratio between market services and nondurable goods. The series in each respective panel are normalized to one for low skilled single men in year 2010. For a comparison with the unfiltered data, see Figure B.3 in Appendix 2.B. Figure B.4 in Appendix 2.B compares the expenditure shares.

To gain some understanding, consider some averages. The average for the ratio of the ratios (the left-hand-side of Equation 2.14) across the years is 1.78, i.e., low skilled men spend significantly more time in home production relative to the capital they employ. To some extent, this can be explained by the fact that the wage rate that low skilled individuals face in the market is lower, which increases their opportunity cost of hours in home production.[10] More specifically, the average over the years for this ratio is 1.34. However, given that capital and labor are estimated to be gross-complements in home production, this can only explain a quarter of the cross sectional dispersion in the ratio of interest.[11] Through the lens of the model, it must be the case that low skilled men are relatively more efficient at using capital as compared to their hours in home production, compared to high skilled men. With these factors of production being gross complements, the household substitutes from capital to labor when the ratio z_k/z_l increases, which can be more easily seen in Equation 2.5.

What about the changes over time? For high skilled men, the ratio in the year 2018 was merely 0.33 times its value in the year 1962. Over the same period, the relative price of capital versus the price of hours in home production (the wage rate) decreased substantially: in the year 2018 the price of capital relative to the price of hours in home production was about 0.135 times its value in the year 1962. Without any changes in the ratio z_k/z_l, the hours-to-capital ratio would have had to fall substantially more. Since capital-augmenting technology grew more than labor-augmenting technology between the years 1962 and 2018, and labor and capital are gross complements in home production, this counteracts the decrease in the relative price and induces less substitution away from hours in home production.

Capital-to-nondurable goods ratio: Figure 2.3b graphs model-generated time series for the capital-to-nondurable goods ratios together with their data counterparts. High skilled men use more capital relative to nondurable goods in home production than

[10]See the development of wage rates in Figure B.2 in Appendix 2.B.

[11]Consider again the averages. Since the ratio on the left-hand-side is 1.78, the ratio of wage rates is 1.34, and $\sigma_{kl} = 0.66$, it can explain about $\frac{1.34^{0.66}-1}{1.78-1} \approx 27\%$ of the level difference between low and high skilled single men.

do low skilled men, and for both of these types of households the ratio has been increasing over time.

The ratio can be decomposed as follows:

$$\frac{K_{jt}}{N_{jt}} = \frac{K_{jt}}{Q_{jt}} \times \frac{Q_{jt}}{N_{jt}} \tag{2.15}$$

First, the ratio is affected by how the household substitutes between capital and the composite Q. Second, the ratio is affected by how the household substitutes between the different factors of production that make up the composite Q. The ratio K/Q is actually lower among high skilled. Considering that the factors of production in the composite Q are capital and labor, and that the ratio K/L_h is higher among high skilled households, this might seem counterintuitive. The factor augmenting technologies z_k and z_l is the reason why K/Q is higher among high skilled. More specifically, it is z_l that explains why K/Q is lower among high skilled men, since this technology is higher among high skilled men and has a negative effect on the ratio. When z_l increases, so does Q, which lowers K/Q.[12]

It must be the case that the ratio Q/N is higher among high skilled than among low skilled. The difference between household types is given by

$$\frac{Q_{high,t}/N_{high,t}}{Q_{low,t}/N_{low,t}} = \left(\frac{P_{q,high,t}}{P_{q,low,t}}\right)^{-\sigma_{nq}} \left(\frac{z_{n,high,t}}{z_{n,low,t}}\right)^{1-\sigma_{nq}} \tag{2.16}$$

where P_q denotes the implicit price of the composite Q, derived in Appendix 2.A.1. Both parts on the right-hand-side of Equation (2.16) contribute to that the ratio on the left-hand-side is greater than one. The implicit price P_q, however, only contributes a small part, while

[12]In Appendix 2.A.2 I show why this is the case.

it is the nondurable augmenting technology ratio that explains most of it.[13]

The capital-to-nondurable goods ratio has increased substantially over time and was in the year 2018 around three times as large as it was in the year 1962, for both types of households. I use the same decomposition as before to identify the factors driving the change over time. Both K/Q and Q/N increased from the year 1962 to the year 2018 and thus contributed to the increase in K/N. However, they did not grow at the same time: while Q/N increased up until the beginning of the 1990s and then did not grow anymore, K/Q started growing at the beginning of the 1990s up until the year 2018. That K/Q was more-or-less constant up until the beginning of the 1990s is explained by that there being two counteracting forces that balanced each other. On the one side, the relative price P_q/P_k grew, which had a positive effect on the ratio K/Q. On the other side, z_k also grew, which affected K/Q in the opposite direction. After the beginning of the 1990s, both P_q/P_k and z_k continued to grow, but this time the first effect dominated the second.[14]

That the ratio Q/N first increased and then changed only marginally from the 1990s and onward was mostly caused by the evolution of z_n. Up until the 1990s, z_n grew rapidly which caused the households to substitute away from N and towards Q, since these factors are gross complements in home production. When z_n stopped growing thereafter, this driver behind substitution stopped.[15]

Market services-to-nondurable goods ratio: Figure 2.3c shows that the market services-to-nondurable goods ratio increased

[13]See Appendix 2.A.3 for more details.
[14]See Appendix 2.A.4 for more details.
[15]See Appendix 2.A.3 for more details.

for many years and then decreased. In a standard framework without home production and homothetic preferences, this is impossible to generate, given that the relative price of nondurable goods decreased monotonically over time. As is also clearly visible in the figure, the ratio is substantially larger among high skilled men, and this is what I start analyzing.

The market services to nondurable goods ratio can be decomposed as follows:

$$\frac{C_{mjt}}{N_{jt}} = \frac{C_{mjt}}{C_{hjt}} \times \frac{C_{hjt}}{N_{jt}} \tag{2.17}$$

First, it depends on how households substitute between market produced and home produced services. Second, it depends on the quantity of nondurable goods used in home production. While C_m and N can be measured in the data, C_h is a model-constructed object. Given the estimated values of the parameters, the ratio C_m/C_h is lower among high skilled. This is driven by the produced quantity of home services being significantly higher among high skilled. The reason for this is that they are more efficient at it, which, even though their market wage rate is higher, leads to a lower implicit price of home production (see Appendix 2.A.5 for more details).

The reason why the market services to nondurable goods ratio is higher among high skilled men than among low skilled men is that in relation to how many units of home services that are produced, high skilled single men use fewer units of nondurable goods. The implicit price of home production also plays a role in determining the ratio C_h/N, but the effect now goes in the opposite direction: a lower implicit price of home production among high skilled increases C_h/N, and hence increases C_m/N. The factor augmenting technology z_n also plays an important role: as N and Q are gross complements in home production, $z_{n,\mathrm{high}} > z_{n,\mathrm{low}}$ leads to more substitution towards Q and away from N for high skilled, which also has a positive effect on the ratio C_m/N.

Let us look more closely at the drivers behind the non-monotonic time-series development of C_m/N. For both types of households, the quantity ratio between market and home produced services decreased up until the early 1990s. Between then and 2018, it decreased but the decrease was small. This is caused by a decrease and then a marginal increase of P_h. The decrease that lasted around three decades caused households to substitute away from market services and into home produced services. Wages did not decrease over this period and therefore the decrease must be driven by factor augmenting technologies and falling prices of capital and nondurable goods. With regards to technology, Figure 2.2 suggests that it is in particular the increases in z_n and z_k during this period that caused P_h to decrease.

An increase in C_h/N explains why C_m/N increased up until the 1990s. That is, in proportion to how much home services that were produced, the quantity of nondurable goods used decreased. This is driven both by the fact that the price ratio P_n/P_h increased during this period and by the increase in z_n. From the early 1990s, the quantity of nondurable goods increased somewhat faster than the quantity of home services, leading the ratio C_h/N to fall, albeit modestly.

In summary: up until the 1990s, households did substitute towards home services which, ceteris paribus, increased the demand for nondurable goods. But at the same time, households substituted away from nondurable goods in home production, to a large extent as a result of that they began using them more efficiently. These two forces counteracted each other, but the effect from substitution away from nondurable goods in home production dominated. From the beginning of the 1990s, the effects have gone in the same direction; households have been substituting somewhat back towards market services and have also become more reliant on nondurable goods in home production. These effects have caused the ratio C_m/N to decrease during the last three decades.

The allocation of time for men

This section focuses on understanding time-use. I do not explicitly focus on the level of expenditures. It has already been shown that the model does well in terms of matching the time series for how expenditures are allocated. Hence the level of expenditures, and the change in these, generated by the model will typically be in line with the data as long as hours worked in the market are in line with the data.

To what extent can the model match the time-use data, and how can it be understood? Figures 2.4b and 2.4c show market hours and hours in home production, respectively. Figure 2.4a shows the sum of these. That total hours worked decreased between years 1962 and 2018 is the case in both the model and the data, although the model-generated decrease is lower. That the model underpredicts the decline for total hours is mainly due to the fact that it underpredicts the decline in market hours. Hours in home production increase for both low and high skilled men in the data between the years 1962 and 2018. Qualitatively, the model matches this fact. Quantitatively, the model-generated increase is similar in the model and the data for low skilled single men, but the model overpredicts the increase somewhat for high skilled single men.

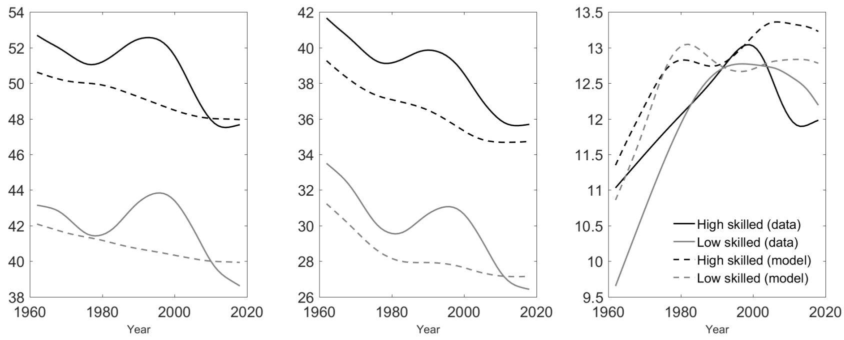

(a) Total *(b) Market* *(c) Home*

Figure 2.4: *Hours worked for men: model and data*

Note: *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production. For a comparison with the unfiltered data, see Figure B.5 in Appendix 2.B.

Total hours worked in the model, for men, are given by

$$L_{jt} = \left(\frac{W_{jt}}{P_{jt}}\right)^{\frac{1-\gamma}{\gamma+\varphi}} \chi_j^{-\frac{1}{\gamma+\varphi}} \tag{2.18}$$

where P is the implicit price of the consumption index $C = \left(C_m^{\frac{\sigma_{mh}-1}{\sigma_{mh}}} + C_h^{\frac{\sigma_{mh}-1}{\sigma_{mh}}}\right)^{\frac{\sigma_{mh}}{\sigma_{mh}-1}}$.[16] In Appendix 2.A.7 I show that the long-run change in $\frac{W}{P}$ is positive as long as the productivity terms z_n, z_l and z_k are growing, which they do via restriction. From Equation (2.18), it is easy to see that a necessary condition to be able to account for the decrease in total hours worked is therefore to have $\gamma > 1$. In Section 2.5.2 I discussed that the implicit home price index is lower for high skilled households. This leads to the index

[16]This implicit price is derived in Appendix 2.A.6.

P being lower and that the ratio $\frac{W}{P}$ will be unambiguously higher for high skilled households. If $\chi_{high} = \chi_{low}$, this would lead to the prediction that high skilled households work fewer hours, which is not what is observed in the data. More specifically, from Equation (2.18) we can learn that $\chi_{high} < \chi_{low}$ is necessary to explain why total hours worked are higher among high skilled single men than among low skilled single men. As was presented in Table 2.3, the estimate for χ_{high}/χ_{low} is 0.47.

Hours worked in home production increased until the late 1990s. The model does a fairly good job in terms of generating predictions that align with the data. One discrepancy that is particularly noticeable, however, is that the increase in the model occurs during the first two decades and thereafter the hours only vary a little. One way of writing the model prediction for hours worked in home production is as a fraction of total hours:

$$L_{hjt} = \eta_{jt} L_{jt} \tag{2.19}$$

where $\eta_{jt} = \left(\frac{W_{jt}/z_{ljt}}{P_{qjt}}\right)^{-\sigma_{kl}} \left(\frac{P_{qjt}}{P_{hjt}}\right)^{-\sigma_{nq}} \left(\frac{P_{hjt}}{P_{jt}}\right)^{-\sigma_{mh}} \frac{W_{jt}/z_{ljt}}{P_{jt}}$. As total working hours fall over time, it is necessary (but not sufficient) that the share of total work hours allocated to home production increases. It is clear from Equation (2.19) that multiple things affect how L_h develops and my results show that these forces are not working in the same direction. The ratios $\frac{W/z_L}{P_q}$ and $\frac{P_q}{P_h}$ grow over time[17] for both household types, which has a negative effect on the hours worked in

[17]While $\frac{W/z_L}{P_q}$ is monotonically increasing over time, $\frac{P_q}{P_h}$ grows until it stabilizes in the early 2000s.

home production. I combine equations (2.18) and (2.19) to get

$$L_{hjt} = \left(\frac{W_{jt}/z_{ljt}}{P_{qjt}}\right)^{-\sigma_{kl}} \left(\frac{P_{qjt}}{P_{hjt}}\right)^{-\sigma_{nq}} \left(\frac{P_{hjt}}{P_{jt}}\right)^{-\sigma_{mh}} \times \left(\frac{W_{jt}}{P_{jt}}\right)^{\frac{1+\varphi}{\gamma+\varphi}} \chi^{-\frac{1}{\gamma+\varphi}} z_{ljt}^{-1} \tag{2.20}$$

which incorporates the effect on L and expresses the development of L_h as a function of parameters and prices. In addition to what was discussed above, which had a negative effect on hours in home production, there are three additional parts that affect how L_h changes over time. First, as z_l grew, the need for L_h diminished. Second, $\frac{P_h}{P}$ became smaller over time and this had a positive effect on home production hours, as households substituted to home services from market services. Third, the real wage $\frac{W}{P}$ increased. This part acts as a level shifter as it increases the overall level of the consumption index.

2.6 Estimation: Women

Model parameters that have been estimated thus far are from here on taken as given. Some additional parameters need to be estimated in order to generate predictions for women. Most of these parameters are of the same type as those that were estimated for men, but are allowed to be different for women. Specifically, these are the household specific levels of the respective factor augmenting technologies, z_l, z_k, and z_n, and the parameter χ that shifts disutility of labor. Two parameters that are completely new are κ_m and κ_h. While the disutility from working in the market and in home production was the same for men, this will not be the case for women. The values for κ_m and κ_h

will be restricted to be the same for both types of women but can change over time. Introducing these two parameters will moreover lead to the implicit prices for women having a different interpretation as compared to men.

2.6.1 Substitution Between Home Hours and Capital

For single women, the ratio between hours in home production and capital is given by

$$\frac{L_{hjt}}{K_{jt}} = \left(\frac{P_{kt}}{W_{jt}}\frac{\kappa_{mt}}{\kappa_{ht}}\right)^{\sigma_{kl}} \left(\frac{z_{kjt}}{z_{ljt}}\right)^{1-\sigma_{kl}} \tag{2.21}$$

The developments of z_k/z_l, in their indexed form, are restricted to be the same as for single men. Given this restriction, the path for κ_m/κ_h in its indexed form is identified by the changes over time in the hours-to-capital ratio. The levels of the ratios κ_m/κ_h and z_k/z_l cannot, however, be separately identified using the hours-to-capital ratio. I assume that $\kappa_m/\kappa_h = 2$ in the year 1962. Given this starting point, I can solve for κ_m/κ_h in each year, and the type-specific levels of z_k/z_l, that jointly minimize the mean squared errors for low and high skilled women for the hours in home production-to-capital ratio.[18] Like for men, part of the cross-sectional variation between women will be explained by differences in wage rates, and part by the difference in the ratio z_k/z_l. The changes over time in the ratio κ_m/κ_h will

[18] The model's ability to hit the hours-to-capital ratio does not depend on the assumed value of κ_m/κ_h in the year 1962, as z_k/z_l can adjust. However, the assumed starting value matters for other moments, like, e.g., the market services-to-nondurable goods ratio, and time-use. A higher starting value for κ_m/κ_h can make the model generate output that better aligns with C_m/N in the data, but will make the prediction for time-use significantly worse. I choose 2 as the starting value because (i) it seems plausible, and (ii) I put significant weight on hitting the time-use data.

partly explain the time-series developments of L_h/K.

2.6.2 Substitution Between Capital and Nondurable Goods

The expression for the capital-to-nondurable goods ratio for women looks like the one for men:

$$\frac{K_{jt}}{N_{jt}} = \Omega_{jt}^{\frac{\sigma_{nq}-\sigma_{kl}}{\sigma_{kl}-1}} \left(\frac{z_{kjt}}{z_{njt}}\right)^{\sigma_{nq}-1} \left(\frac{P_{nt}}{P_{kt}}\right)^{\sigma_{nq}} \tag{2.22}$$

where

$$\Omega_{jt} = 1 + \left(\frac{z_{ljt}P_{kt}}{z_{kjt}W_{jt}}\frac{\kappa_{mt}}{\kappa_{ht}}\right)^{\sigma_{kl}-1} \tag{2.23}$$

i.e., here the ratio κ_m/κ_h shows up. At this stage, only levels for z_k/z_n for each type of woman are pinned down. Values for all other variables or parameters in Equation (2.22) are either known or already fixed.

2.6.3 Substitution Between Market Services and Nondurable Goods

The equation for C_m/N also looks like it does for men:

$$\frac{C_{mjt}}{N_{jt}} = \Gamma_{jt}^{\frac{\sigma_{mh}-\sigma_{nq}}{1-\sigma_{nq}}} P_{nt}^{\sigma_{mh}} z_{njt}^{1-\sigma_{mh}} \tag{2.24}$$

but for women the ratio κ_m/κ_h enters through Γ (that depends on Ω, see Equation (2.9)). Equation (2.24) is used to estimate levels for z_n for each woman. The levels for the factor augmenting technologies z_l and z_k are pinned down once z_n is fixed.

At this point, the ratio κ_m/κ_h is fixed for each year, but the levels for each of the two components that make up the ratio are now estimated, together with χ for each respective type of woman. I target hours worked in the market and in home production for both types of households, in each year, and the estimated values for the parameters minimize the sum of the squared errors.

2.7 Results: Women

2.7.1 Parameters

The estimated cross-sectional differences in the factor augmenting technologies are presented in Table 2.4. Clearly, they imply rather large differences across the two types of women in terms of how productive they are at using each respective factor of production. While high skilled women are significantly more productive at using non-durable goods and their hours, low skilled women are about twice as productive as high skilled women at using capital. Qualitatively, the results are similar to the results for men.

Table 2.4: *Women: estimated cross-sectional differences in factor efficiency*

$\frac{z_{n,high}}{z_{n,low}}$	$\frac{z_{k,high}}{z_{k,low}}$	$\frac{z_{l,high}}{z_{l,low}}$
2.38	0.50	5.16

The estimates for κ_m and κ_h tell the story that market work became less unpleasant while domestic work became more unpleasant over time. This is shown in Figure 2.5a. Compared to in the year 1962, κ_m fell by around 35 percent until the year 2018 while κ_h increased by 123 percent over the same period.[19] Figure 2.5b plots the ratio of

[19] In particular the change in κ_h is substantial. What could the increase capture?

the two and shows a substantial and monotonic decline in the ratio up until the early 2000s, after which the ratio has been fairly constant.

Finally, regarding the general distaste for work, it is estimated that $\chi_{low}/\chi_{high} = 2.28$.

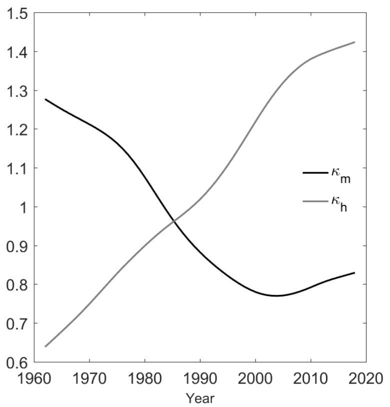

(a) κ_m *and* κ_h

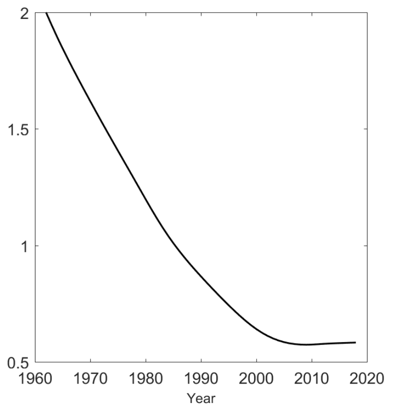

(b) Ratio κ_m/κ_h

Figure 2.5: *Development of social norms*

Note: *Panel (a)*: Estimated values for κ_m and κ_h. *Panel (b)*: Ratio of the estimated values for κ_m and κ_h. Note that the ratio is restricted to start at 2 in year the 1962, see the text for more details.

2.7.2 Model Implied Time Series

Ratios and expenditure shares for women

Hours in home production-to-capital ratio: Having κ_m and κ_h in the model for women lets it produce time series for the hours in home production-to-capital ratio that align well with data (see Figure

---

If it reflects changes in social norms, the increase could be interpreted as a *positive* view on women as home producers decreasing over time.

2.6a). The ratio is substantially lower among high skilled women as compared to low skilled women. In any given period t, the difference between the two types of women is given by

$$\frac{L_{h,high,t}/K_{high,t}}{L_{h,low,t}/K_{low,t}} = \left(\frac{W_{high,t}}{W_{low,t}}\right)^{-\sigma_{kl}} \left(\frac{z_{k,high,t}/z_{l,high,t}}{z_{k,low,t}/z_{l,low,t}}\right)^{1-\sigma_{kl}} \tag{2.25}$$

On average over all the years, the ratio has been 0.36. About a third of this difference is due to the fact that high skilled women met higher wage rates in the market compared to low skilled women, thus increasing the opportunity cost of using time for home production.[20] The more important driver behind the substantial difference is differences in factor augmenting productivities. Specifically, the ratio between how efficiently capital is used relative to how efficiently hours are used is higher among low skilled women than among high skilled women. With capital and labor being gross complements, this induces low skilled women to use relatively more hours relative to capital in home production.

To analyze the drivers of how the hours-to-capital ratio developed over time, I focus on high skilled women. As the development was quantitatively similar for low skilled women, the conclusions apply to them as well. When men were analyzed in Section 2.5.2, it was established that the development of z_k/z_l in isolation had a positive effect on the hours-to-capital ratio. For high skilled single women in the year 2018, the ratio L_h/K is only one fifth of what it was in the year 1962. This decline is substantially larger than for high

[20]Based on the average wage rates for each type of women households, $(W_{high}/W_{low})^{-\sigma_{kl}} = 0.82$. I get to approximately one third by computing $(1-0.82)/(1-0.36) \approx 0.28$

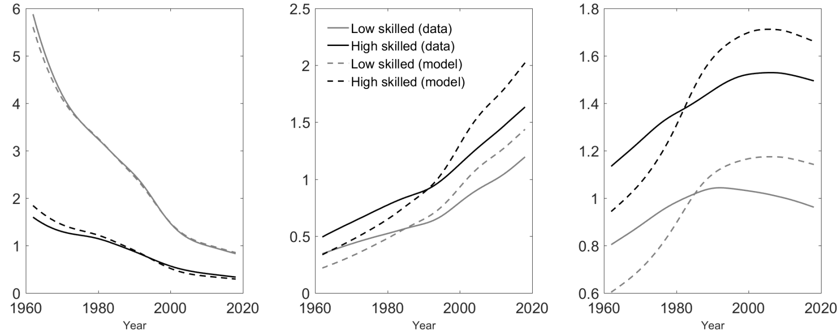

(a) L_h/K *(b)* K/N *(c)* C_m/N

Figure 2.6: *Quantity ratios for women: model and data*

Note: *Panel (a)*: Ratio between hours in home production and capital. *Panel (b)*: Ratio between capital and nondurable goods. *Panel (c)*: Ratio between market services and nondurable goods. The series in each respective panel are normalized to one for low skilled single women in the year 2010. For a comparison with the unfiltered data, see Figure B.6 in Appendix 2.B. Figure B.7 in Appendix 2.B compares the expenditure shares.

skilled single men, for whom the ratio dropped to one third. Much of the decline is explained by the fact that the relative price between capital and hours in home production decreased over time, and the decrease was larger for high skilled women compared to high skilled men (driven by a larger increase in the wage rate for women). Given an elasticity of substitution between hours and capital of 0.66, the greater decline in the relative price between capital and hours was still not enough to explain a larger decline in the hours-to-capital ratio for women. Through the lens of the model, the distaste for market work relative to domestic work must have decreased and quantitatively this is important. When the distaste for market work relative to domestic work decreases, this affects the ratio via two channels. First, in utility terms it becomes more costly to produce home services with hours in home production, leading the household to decrease them relative to capital. Second, it becomes relatively more attractive, in utility terms, to work in the market. This increases market hours, increases the budget, and has a positive effect on capital expenditures.

Capital-to-nondurable goods ratio: Figure 2.6b shows the model-generated time series for the capital-to-nondurable goods ratio together with their data counterparts. As the growth over time for z_k/z_n was generated by targeting the ratios for single men, it is not surprising that the differences between model generated series and data are larger for women. The model still does a good job at predicting the developments, even though it overpredicts the increases somewhat.

Based on the series that are produced by the model, the capital-to-nondurable goods ratio has, on average across all the years, been 1.4 times higher among high skilled women compared to low skilled women. Moreover, the capital-to-nondurable goods ratio has increased substantially over time. I use the same decomposition as for men (see Equation (2.15)) to understand the cross-sectional differences and the development over time.

The conclusions are very similar to those drawn when analyzing cross-sectional differences in the corresponding ratios for men. The ratio K/Q is lower among high skilled women compared to low skilled women, as a result of z_l being higher for the prior type. I refer the reader to the same Appendix 2.A.2 like I did when analyzing men to understand why this must the case. The fact that the capital-to-nondurable goods ratio is higher among high skilled single women is thus driven by the ratio Q/N being higher. The main reason why the ratio Q/N is higher, and the ratio K/N is higher, among high skilled women than among low skilled women, is that high skilled women use nondurable goods more efficiently and, since the factors of production are gross complements, substitute towards the composite Q, which is partly made up of capital. I go through this in more detail in Appendix 2.A.9.

The increase in the capital-to-nondurable goods ratio between the years 1962 and 2018 is explained by both ratios K/Q and Q/N increasing. But like for men, it was Q/N that was the dominant force behind the increase up until the early 1990s, from which point in time K/Q took over as the driving force.

Up until the 1990s, both capital-augmenting technology and nondurable goods-augmenting technology grew rapidly, at roughly equal rates. Quantitatively, it was, however, primarily the growth in z_n that had an impact on Q/N, as a result of the low elasticity of substitution between N and Q. Although z_k kept growing after the 1990s, z_n did not and the ratio Q/N plateaued. At the beginning of the 1990s, the ratio K/Q started growing faster, and took over as the driver behind the increase in the capital-to-nondurable goods ratio. Two developments were roughly equally important behind this increase. First, the relative price between hours and capital (W/P_k) increased, thus increasing the demand for capital. Second, the ratio κ_h/κ_m increased, also making the women substitute away from hours in home production and towards market work and capital. Absent the increases in any of these two ratios, there would not have been an increase in the K/Q ratio (see Appendix 2.A.9 in which I show this through counterfactual scenarios).

Market services-to-nondurable goods ratio: I now turn to analyzing the market services-to-nondurable goods ratio. In Appendix 2.A.10 I go into details about why the ratio is higher among high skilled women than among low skilled women. The most important factor is that low skilled women use nondurable goods less efficiently. This has two opposing effects on the ratio. First, that z_n is lower among low skilled women leads to their implicit price of home services being larger, thus leading them to consume more market services relative to home produced services compared to high skilled women. But, second, a lower z_n among low skilled women also leads them to use relatively more nondurable goods in home production since the composite Q and nondurable goods are gross complements. Out of these two counteracting forces, the latter effect dominates and it can be concluded that the market services-to-nondurable ratio is higher among high skilled women as a result of them, relative to low skilled women, using less nondurable goods in home production.

The hump-shaped development for the ratio between quantities of market services and nondurable goods is, like for men, present for women (see Figure 2.6c). The ratio can be better understood by using the decomposition from Equation (2.17), which says that it depends on how households substituted between market and home services, and how much nondurable goods that are used to produce the latter. I again refer the reader to Appendix 2.A.10 for a more in-depth analysis.

An important margin along which women substituted, that contributed to the increase in the ratio before the year 2009, was from home produced services to market services. This differs vastly from what men did, who substituted in the other direction. This development was due to shifts in the relative distastes of market versus domestic work. More specifically, a decreasing relative distaste for market work increased the implicit price of home services and this, in turn, led women to substitute towards market produced services. The ratio κ_h/κ_m stabilized after the year 2009. Two other components, more specifically the prices per efficiency units of hours in home production and capital, respectively, fell after the year 2009. These developments led the implicit price of home services to fall and women to some degree to substitute away from market produced services after 2009.

C_h/N increased up until the year 1992 as a result of nondurable goods being used more efficiently over time, causing households to substitute away from them and into the composite Q, since Q and N are gross complements. This increase in z_n stopped after the year 1992, while the relative price P_n/P_h continued to decrease like it also had prior to this point in time. The latter decrease had an effect on the ratio in question as nondurable goods became more affordable as an input in home production, but still, because of the low value of σ_{nq}, only had a small quantitative effect.

In summary, two things stand out in terms of how important they were for explaining the development of the market services-to-nondurable goods ratio. First, nondurable goods augmenting technology grew rapidly up until 1992, causing women to substitute away from nondurable goods in home production, and this contributed significantly to the increase up to that point. After 1992, z_n fell and had a negative effect on the ratio. Over the period 1962-2009, women continuously increased their consumption of market produced services relative to home produced services, mainly as a result of the decreasing relative distaste for market work (this is the second thing that stands out as important). The fact that women increasingly substituted away from home produced services also during the years 1992-2009 dominated the decrease in C_h/N over the same period and hence the ratio C_m/N increased. When the relative distaste for market work stopped decreasing and other prices and technology developments all led home services to become relatively cheaper, households started substituting towards home services and away from market services. Thus, after 2009, both C_h/N and C_m/C_h fell, two factors joining forces in pushing down C_m/N.

The allocation of time among women

The three panels of Figure 2.7 illustrate the development of total hours worked, hours worked in the market, and hours worked in home production, respectively, in the data and model. Overall, the model aligns well with the data, which, given that the paths for all the productivity series as well as the path for κ_h/κ_m are not targeting these series specifically, is quite impressive. Compared to the strong trends in time-use for married women (which are not included in this paper), the trends over time for single women are mostly modest. The amount of hours that high skilled single women supplied to the market did increase until the late 1990s, by around four hours, but then fell back to the levels at which they started. Over the period as a whole, hours worked in home production decreased among high skilled women, but only by about an hour. While market hours among low skilled women varied only little over time and showed no trend, the hours worked in home production clearly decreased, by eight hours.

Getting to the bottom of what is contributing to cross-sectional differences and changes over time in time-use is challenging. I will summarize my findings in this section, but for more details I refer the reader to Appendix 2.A.11, in which I rely heavily on counterfactual scenarios to identify the drivers.

Differences in wage rates have significant effects on the level differences between the two types of women. However, equalizing wages

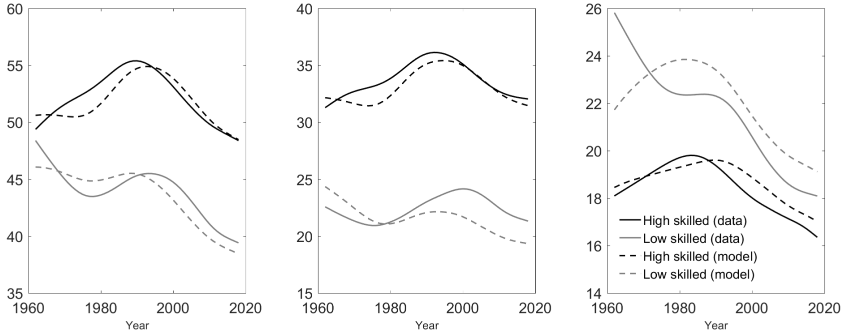

(a) Total *(b) Market* *(c) Home*

Figure 2.7: *Hours worked for women: model and data*

Note: *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production. For a comparison with the unfiltered data, see Figure B.8 in Appendix 2.B.

across the two groups would lead to larger differences. The changes in wage rates over time only had small effects on hours worked in home production. Up until the 1990s, the effect is negligible also on market hours for high skilled women, but thereafter the increase in the wage rate had a dampening effect. With respect to the changes over time in market hours for low skilled women, the exercise suggests that they would have been almost 10 percent higher in 2018 had the wages not grown.

Cross-sectional differences in the levels of each factor augmenting technology had considerable effects on market hours. The level differences also affected the time trends, a consequence of the non-linear nature of a CES structure, like the one employed in the current paper. In particular, these non-linearities result in that the effects from changes in variables or parameters such as, e.g., the wage rate or social norms, depend on the levels of the factor-augmenting technologies. Interestingly, the market hours for high skilled decrease over time and the market hours for low skilled increase in the counterfactual scenario where I give low skilled women the factor augmenting technology levels from high skilled women, and vice versa.

Changes in social norms are crucial for understanding how women allocated their time. Without any changes in social norms, high skilled women would have significantly decreased hours of market work. High skilled women would have worked about four hours more in home production in the year 2018 compared to in the year 1962, and the change would amount to six hours for low skilled women. This is clearly in contrast to the actual development.

2.8 A Closer Look at the Importance of Social Norms

The importance of social norms has to some extent already been analyzed. In Section 2.7.2, it was, e.g., found that social norms are crucial for explaining the fall in domestic hours among women over time. Moreover, the hours that high skilled women worked in the market would have fallen considerably over time had not the distaste for market work decreased relative to the distaste for domestic work. The current section complements what has already been done, by looking at how the three ratios C_m/N, K/N, and L_h/K would have developed without any changes in social norms.

Figure 2.8 illustrates the results. In the right-hand panel, each series is indexed to its starting value in the year 1962. The results highlight that there is a tension between matching the development of the hours worked at home-to-capital ratio (L_h/K) and the developments of hours worked in the market and in home production, or hitting the capital-to-nondurable goods ratio and the market services-to-nondurable goods ratio. Shutting down the developments of social norms gives a model output that is more in line with the data for K/N and C_m/N, but less in line in terms of L_h/K. We know from before that the model-generated time series for time-use do not even resemble their data counterparts with unchanged norms.

With that said, I still analyze how the changes over time in each of these ratios would have been different had social norms not changed. First, the decrease in how many hours of work that go into home production relative to capital would have decreased significantly less. This was already discussed when the importance of social norms was evaluated in Section 2.7.2. From Figures 2.8c and 2.8d it can be seen that the model-generated series for K/N are more in line with the data when social norms do not change. This is especially clear from Figure 2.8d. The reason for the slower increase is a significantly smaller increase in the ratio K/Q. This, in turn, is explained by households not substituting away from hours in home production and into capital as much, since the relative distaste for market work does not fall. Lastly, the model results imply that the increase in the market services-to-nondurable goods ratio would have been more modest with unchanged social norms. A consequence of having constant social norms is that the implicit price of home production is trending down over time, as opposed to up like it is estimated to do. The decreasing implicit price of home production would have caused households to substitute from market produced services to home produced services, which is why the ratio increases by so much less, and even starts decreasing earlier, in the counterfactual scenario.

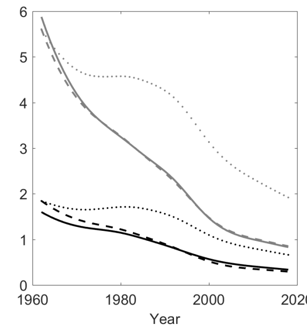

(a) $\mathsf{L_h/K}$

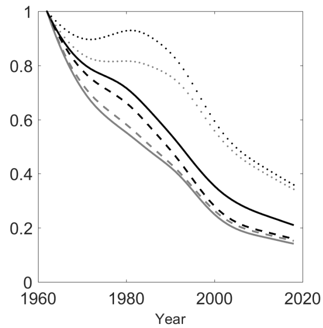

(b) $\mathsf{L_h/K}$ $(1962 = 1)$

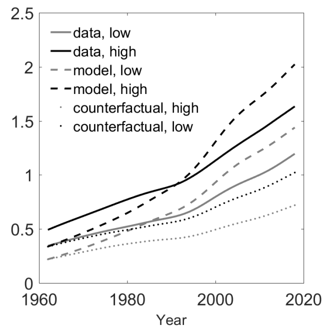

(c) $\mathsf{K/N}$

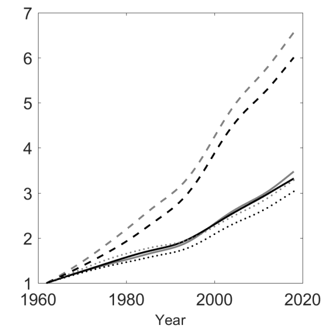

(d) $\mathsf{K/N}$ $(1962 = 1)$

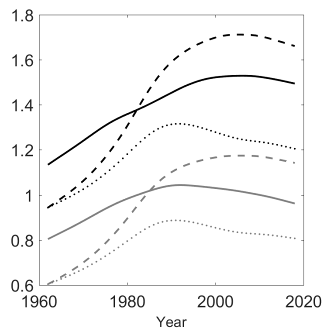

(e) $\mathsf{C_m/N}$

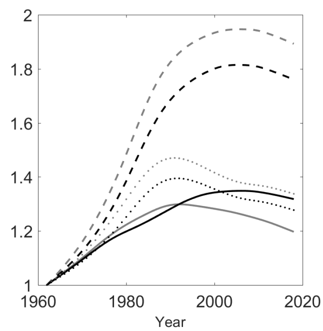

(f) $\mathsf{C_m/N}$ $(1962 = 1)$

Figure 2.8: *Ratios: women, with constant social norm*

Note: *Panel (a)*: Ratio between hours in home production and capital. *Panel (c)*: Ratio between capital and nondurable goods. *Panel (d)*: Ratio between market services and nondurable goods. The series in panels (a), (b), and (c) are normalized to one for low skilled women in the year 2010. Panels (b), (d), and (f) index series from the panel to the left on the same row and start at one in the year 1962.

2.9 Conclusions

The current paper proposes a unified analysis of expenditures and time-use, via a framework in which households can produce services at home. It is shown that the model, which uses stable homothetic preferences and standard functional forms for home production, can match data about expenditures and time-use, both in the cross-section and the developments over time. In particular, it can explain why the ratio between quantities of market services and nondurable goods increased over many years, despite an increase in the relative price, which most models need non-homothetic preferences to generate.

An ex-ante plausible story behind why the market services-to-nondurable goods ratio increased over many decades is that home services became marketized. I find this not to be the case for men. Instead, I find that men increased the consumption of home services relative to market services. The driver behind this was that the implicit price of home production fell. This, in turn, was partly explained by the decreasing relative prices of capital and nondurable goods (which are used as factors of production at home), and partly by home production becoming more efficient. Through the lens of the model, the increase in the market services-to-nondurable goods ratio resulted from a substantial growth in nondurable goods-augmenting technology up until the 1990s, which caused men to substitute away from nondurable goods in home production, as production factors are estimated to be gross compliments.

Women did also increase the market services-to-nondurable goods ratio for many decades. Like men, they did also substitute away from nondurable goods in home production. But contrasting the development for men, women did also reduce their consumption of home produced services relative to market services. Changes in social norms lie behind this development: a decrease in the relative distaste for market work relative to domestic work led to women working more in the market and less in home production. Absent changes in social norms, the developments for women would have been vastly different, both in terms of how they allocate their time, and in terms of how expenditures were allocated.

The current paper does not model couple households. Extending the model in this direction could be straightforward. The most obvious alternative would be to add yet another nest in the home production function, that combines two spouses' home production hours into a composite, and that this composite, in the next nest is then combined with capital. Social norms are most likely even more important for explaining changes in market hours for married women. Considering the dramatic rise in labor supply among married women over many decades, I think that such an analysis would be very interesting. Moreover, with a structure similar to the proposed one, the development of men's time-use would also be affected by these changes in social norms.

References

Acemoglu, D. and Guerrieri, V. (2008). Capital deepening and nonbalanced economic growth. *Journal of political Economy*, 116(3):467–498.

Aguiar, M., Hurst, E., and Karabarbounis, L. (2013). Time use during the great recession. *American Economic Review*, 103(5):1664–1696.

Almgren, M. (2023). The allocation of expenditures and time over time. Technical report, IIES.

Aruoba, S. B., Davis, M. A., and Wright, R. (2016). Homework in monetary economics: Inflation, home production, and the production of homes. *Review of Economic Dynamics*, 21:105–124.

Balleer, A., Gomez-Salvador, R., and Turunen, J. (2014). Labour force participation across europe: A cohort-based analysis. *Empirical Economics*, 46:1385–1415.

Baxter, M. and Jermann, U. J. (1999). Household production and the excess sensitivity of consumption to current income. *American Economic Review*, 89(4):902–920.

Becker, G. S. (1965). A theory of the allocation of time. *The economic journal*, 75(299):493–517.

Benhabib, J., Rogerson, R., and Wright, R. (1991). Homework in macroeconomics: Household production and aggregate fluctuations. *Journal of Political economy*, 99(6):1166–1187.

Boerma, J. and Karabarbounis, L. (2021). Inferring inequality with home production. *Econometrica*, 89(5):2517–2556.

Boppart, T. (2014). Structural change and the kaldor facts in a growth model with relative price effects and non-gorman preferences. *Econometrica*, 82(6):2167–2196.

Boppart, T. and Ngai, L. R. (2021). Rising inequality and trends in leisure. *Journal of Economic Growth*, 26(2):153–185.

Bridgman, B. (2016). Home productivity. *Journal of Economic Dynamics and control*, 71:60–76.

Buera, F. J. and Kaboski, J. P. (2012a). The rise of the service economy. *American Economic Review*, 102(6):2540–69.

Buera, F. J. and Kaboski, J. P. (2012b). Scale and the origins of structural change. *Journal of Economic Theory*, 147(2):684–712.

Burda, M. C. and Hamermesh, D. S. (2010). Unemployment, market work and household production. *Economics Letters*, 107(2):131–133.

Calvo, P. A., Lindenlaub, I., and Reynoso, A. (2021). Marriage market and labor market sorting. Technical report, National Bureau of Economic Research.

Caselli, F. and Coleman II, W. J. (2001). The us structural transformation and regional convergence: A reinterpretation. *Journal of political Economy*, 109(3):584–616.

Comin, D., Lashkari, D., and Mestieri, M. (2021). Structural change with long-run income and price effects. *Econometrica*, 89(1):311–374.

Cravino, J., Levchenko, A., and Rojas, M. (2022). Population aging and structural transformation. *American Economic Journal: Macroeconomics*, 14(4):479–498.

Duernecker, G. and Herrendorf, B. (2018). On the allocation of time–a quantitative analysis of the roles of taxes and productivities. *European Economic Review*, 102:169–187.

Eckstein, Z. and Lifshitz, O. (2011). Dynamic female labor supply. *Econometrica*, 79(6):1675–1726.

Einarsson, T. and Marquis, M. H. (1997). Home production with endogenous growth. *Journal of Monetary Economics*, 39(3):551–569.

Fang, L. and Zhu, G. (2017). Time allocation and home production technology. *Journal of Economic Dynamics and Control*, 78:88–101.

Fernández, R. (2013). Cultural change as learning: The evolution of female labor force participation over a century. *American Economic Review*, 103(1):472–500.

Gomme, P., Kydland, F. E., and Rupert, P. (2001). Home production meets time to build. *Journal of Political Economy*, 109(5):1115–1131.

Gomme, P., Rogerson, R., Rupert, P., and Wright, R. (2004). The business cycle and the life cycle. *NBER macroeconomics annual*, 19:415–461.

Greenwood, J., Guner, N., Kocharkov, G., and Santos, C. (2016). Technology and the changing family: A unified model of marriage, divorce, educational attainment, and married female labor-force participation. *American Economic Journal: Macroeconomics*, 8(1):1–41.

Greenwood, J. and Hercowitz, Z. (1991). The allocation of capital and time over the business cycle. *Journal of political Economy*, 99(6):1188–1214.

Greenwood, J., Rogerson, R., and Wright, R. (1995). Household production in real business cycle theory in t. cooley (ed.), frontiers of business cycle research.

Greenwood, J., Seshadri, A., and Yorukoglu, M. (2005). Engines of liberation. *The Review of Economic Studies*, 72(1):109–133.

Hassler, J., Krusell, P., and Olovsson, C. (2021). Directed technical change as a response to natural resource scarcity. *Journal of Political Economy*, 129(11):3039–3072.

Herrendorf, B., Rogerson, R., and Valentinyi, A. (2013). Two perspectives on preferences and structural transformation. *American Economic Review*, 103(7):2752–2789.

Herrendorf, B., Rogerson, R., and Valentinyi, A. (2014). Growth and structural transformation. *Handbook of economic growth*, 2:855–941.

Hubmer, J. (2023). The race between preferences and technology. *Econometrica*, 91(1):227–261.

Jones, L. E., Manuelli, R. E., and McGrattan, E. R. (2015). Why are married women working so much? *Journal of Demographic Economics*, 81(1):75–114.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1):35–78.

Kongsamut, P., Rebelo, S., and Xie, D. (2001). Beyond balanced growth. *The Review of Economic Studies*, 68(4):869–882.

Lee, J. E. (2021). Marriage and misallocation: Evidence from 70 years of us history.

Leukhina, O. and Yu, Z. (2021). Home production and leisure during the covid-19 recession. *The BE Journal of Macroeconomics*, 22(1):269–306.

McGrattan, E. R., Rogerson, R., and Wright, R. (1997). An equilibrium model of the business cycle with household production and fiscal policy. *International Economic Review*, pages 267–290.

Moro, A., Moslehi, S., and Tanaka, S. (2017). Does home production drive structural transformation? *American Economic Journal Macroeconomics*, pages 116–146.

Ngai, L. R. and Petrongolo, B. (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, 9(4):1–44.

Ngai, L. R. and Pissarides, C. A. (2007). Structural change in a multisector model of growth. *American economic review*, 97(1):429–443.

Ngai, L. R. and Pissarides, C. A. (2008). Trends in hours and economic growth. *Review of Economic dynamics*, 11(2):239–256.

Olovsson, C. (2009). Why do europeans work so little? *International Economic Review*, 50(1):39–61.

Olovsson, C. (2015). Optimal taxation with home production. *Journal of Monetary Economics*, 70:39–50.

Ramey, V. (2007). How much has leisure really increased since 1965? *University of California at San Diego Working Paper*.

Rogerson, R. (2008). Structural transformation and the deterioration of european labor market outcomes. *Journal of political Economy*, 116(2):235–259.

Rupert, P., Rogerson, R., and Wright, R. (2000). Homework in labor economics: Household production and intertemporal substitution. *Journal of Monetary Economics*, 46(3):557–579.

Appendices

2.A Appendix

2.A.1 Implicit price of Q

The cost of producing Q is

$$\mathsf{cost}_{qjt} = P_{kt}K_{jt} + W_{jt}L_{hjt} \tag{2.26}$$

Use Equation (2.5) to substitute for L_h in Equation (2.26):

$$\mathsf{cost}_{qjt} = \Omega_{jt}P_{kt}K_{jt} \tag{2.27}$$

Use Equation (2.7) to substitute for K in Equation (2.27) and get

$$\mathsf{cost}_{qjt} = \Omega_{jt}^{\frac{1}{1-\sigma_{kl}}} \frac{P_{kt}}{z_{kjt}} Q_{jt} \tag{2.28}$$

Hence the implicit price is

$$P_{qjt} = \Omega_{jt}^{\frac{1}{1-\sigma_{kl}}} \frac{P_{kjt}}{z_{kjt}} \tag{2.29}$$

Using the definition of Ω from Equation (2.6) it can also be written as

$$P_{qjt} = \left[\left(\frac{P_{kt}}{z_{kjt}} \right)^{1-\sigma_{kl}} + \left(\frac{W_{jt}}{z_{ljt}} \right)^{1-\sigma_{kl}} \right]^{\frac{1}{1-\sigma_{kl}}} \tag{2.30}$$

2.A.2 The Ratio K/Q for Men and Women

Starting from how the composite Q_{jt} is produced:

$$Q_{jt} = \left((z_{kjt} K_{jt})^{\frac{\sigma_{kl}-1}{\sigma_{kl}}} + (z_{ljt} L_{hjt})^{\frac{\sigma_{kl}-1}{\sigma_{kl}}} \right)^{\frac{\sigma_{kl}}{\sigma_{kl}-1}} \tag{2.31}$$

one can easily get to the following:

$$\frac{K_{jt}}{Q_{jt}} = \left[z_{kjt}^{\frac{\sigma_{kl}-1}{\sigma_{kl}}} + z_{ljt}^{\frac{\sigma_{kl}-1}{\sigma_{kl}}} \left(\frac{K_{jt}}{L_{hjt}} \right)^{-\frac{\sigma_{kl}-1}{\sigma_{kl}}} \right]^{-\frac{\sigma_{kl}}{\sigma_{kl}-1}} \tag{2.32}$$

The partial derivatives with respect to z_{kjt} and z_{ljt} are both negative, while the partial derivative with respect to K_{jt}/L_{hjt} is positive. The ratio K_{jt}/L_{hjt} is higher among high skilled men and women, and z_{kjt} is lower. These two parts increase K_{jt}/Q_{jt}. Thus the reason for why K_{jt}/Q_{jt} is lower for high skilled men and women must come from z_{ljt} being higher.

2.A.3 The Ratio Q/N for Men

Combining Equation (2.6) with Equation (2.29) yields

$$\frac{Q_{jt}}{N_{jt}} = \left(\frac{P_{nt}}{P_{qjt}} \right)^{-\sigma_{nq}} z_{njt}^{1-\sigma_{nq}} \tag{2.33}$$

Taking the ratio between households:

$$\frac{Q_{high,t}/N_{high,t}}{Q_{low,t}/N_{low,t}} = \left(\frac{P_{q,high,t}}{P_{q,low,t}} \right)^{-\sigma_{nq}} \left(\frac{z_{n,high,t}}{z_{n,low,t}} \right)^{1-\sigma_{nq}} \tag{2.34}$$

On average over all the years, the ratio on the left-hand-side is 1.94, the first part on the right-hand-side (the ratio raised to its power) is 1.06 and the second part on the right-hand-side (the ratio raised to its power) is 1.83. The cross-sectional difference in the ratio is thus primarily explained by the differences in z_n, and to only a small extent by the differences in P_q.

The ratio Q/N increased up until the early 1990s and thereafter changed only marginally. The relative price P_n/P_q has been trending downwards over the whole period for both types, which had a positive effect on Q/N. The important thing that changed at the beginning of the 1990s, which caused Q/N to stop growing, was that z_n stopped increasing. Since Q and N are gross substitutes in home production, an increase in z_n leads to the household substituting towards Q and away from N, which was something that increases in z_n caused households to do prior to the early 1990s. When the growth in z_n stopped, so did this substitution.

2.A.4 The Ratio K/Q

This can be written as

$$\frac{K_{jt}}{Q_{jt}} = \left(\frac{P_{qt}}{P_{kt}/z_{kjt}}\right)^{\sigma_{kl}} z_{kjt}^{-1} \tag{2.35}$$

The first part on the right-hand-side contains the relative price between efficiency units of capital and the composite Q. Multiply by z_k on both sides and get

$$\frac{z_{kjt}K_{jt}}{Q_{jt}} = \left(\frac{P_{qjt}}{P_{kt}/z_{kjt}}\right)^{\sigma_{kl}} \tag{2.36}$$

Now the ratio measures the number of efficiency units of capital that is used relative to the value of Q, and is a function of the relative price. The relative price has increased, driven by that the relative price between efficiency units of capital and efficiency units of hours worked in home production has decreased, which can be seen here:

$$\frac{P_{qjt}}{P_{kt}/z_{kjt}} = \Omega_{jt}^{\frac{1}{1-\sigma_{kl}}} \tag{2.37}$$

where $\Omega^{\frac{1}{1-\sigma_{kl}}}$ can be written as

$$\Omega_{jt}^{\frac{1}{1-\sigma_{kl}}} = \left[1 + \left(\frac{W_{jt}/z_{ljt}}{P_{kt}/z_{kjt}}\right)^{1-\sigma_{kl}}\right]^{\frac{1}{1-\sigma_{kl}}} \tag{2.38}$$

An important driver behind this development is the evolution of z_k/z_L, which has more than doubled between the year 1962 and the year 2018. Since capital and hours worked in home production are gross complements in home production, an increase in z_k/z_L has a negative effect on the ratio K/L_h, and in the next step a negative effect on the ratio K/Q.

2.A.5 The Implicit Price of Home Services and C_m/N for Men

Using Equation (2.6), I can write Q as a function of N:

$$Q_{jt} = z_{njt}^{1-\sigma_{nq}} \Omega_{jt}^{\frac{\sigma_{nq}}{\sigma_{kl}-1}} P_{kt}^{-\sigma_{nq}} z_{kjt}^{\sigma_{nq}} P_{nt}^{\sigma_{nq}} N_{jt} \tag{2.39}$$

So the total cost for producing C_h is

$$\text{cost}_{C_{hjt}} = \left[P_{nt} + P_{qjt} z_{njt}^{1-\sigma_{nq}} \Omega_{jt}^{\frac{\sigma_{nq}}{\sigma_{kl}-1}} P_{kt}^{-\sigma_{nq}} z_{kjt}^{\sigma_{nq}} P_{nt}^{\sigma_{nq}}\right] N_{jt} \tag{2.40}$$

Now substitute for P_q using Equation (2.29) and factorize P_n:

$$\mathrm{cost}_{C_{hjt}} = P_{nt}\left[1 + \Omega_{jt}^{\frac{\sigma_{nq}-1}{\sigma_{kl}-1}}\left(\frac{z_{kjt}}{z_{njt}}\frac{P_{nt}}{P_{kt}}\right)^{\sigma_{nq}-1}\right]N_{jt} \tag{2.41a}$$

$$= \Gamma_{jt}P_{nt}N_{jt} \tag{2.41b}$$

where I used the definition for Γ from Equation (2.9). Substitute for N using Equation (2.10) and get

$$\mathrm{cost}_{C_{hjt}} = \Gamma_{jt}^{\frac{1}{1-\sigma_{nq}}}\frac{P_{nt}}{z_{njt}}C_{hjt} \tag{2.42}$$

Hence, the implicit price index for home services is

$$P_{hjt} = \Gamma_{jt}^{\frac{1}{1-\sigma_{nq}}}\frac{P_{nt}}{z_{njt}} \tag{2.43}$$

Let us look closer at this price index. Using the definitions of both Γ and Ω gives

$$P_{hjt} = \left[1 + \left\{1 + \left(\frac{z_{ljt}P_{kt}}{z_{kjt}W_{jt}}\right)^{\sigma_{kl}-1}\right\}^{\frac{\sigma_{nq}-1}{\sigma_{kl}-1}}\left(\frac{z_{kjt}P_{nt}}{z_{njt}P_{kt}}\right)^{\sigma_{nq}-1}\right]^{\frac{1}{1-\sigma_{nq}}}\frac{P_{nt}}{z_{njt}} \tag{2.44}$$

After som algebra, this can instead be written as

$$P_{hjt} = \left[\left(\frac{z_{njt}}{P_{nt}}\right)^{\sigma_{nq}-1} + \left\{\left(\frac{z_{kjt}}{P_{kt}}\right)^{\sigma_{kl}-1} + \left(\frac{z_{ljt}}{W_{jt}}\right)^{\sigma_{kl}-1}\right\}^{\frac{\sigma_{nq}-1}{\sigma_{kl}-1}}\right]^{\frac{1}{1-\sigma_{nq}}} \tag{2.45}$$

or using the expression for P_q from Equation (2.30):

$$P_{hjt} = \left[\left(\frac{P_{nt}}{z_{njt}} \right)^{1-\sigma_{nq}} + P_{qjt}^{1-\sigma_{nq}} \right]^{\frac{1}{1-\sigma_{nq}}} \tag{2.46}$$

From Equation (2.45) we can learn that $dP_h/dW > 0$, and from Equation (2.46) that $dP_h/dz_i < 0$, $\forall i \in \{n, k, l\}$. In the model, the ratio between the quantity of market services and household services is given by

$$\frac{C_{mjt}}{C_{hjt}} = P_{hjt}^{\sigma_{mh}} \tag{2.47}$$

Given observed quantities and the estimated parameter values, the ratio is lower for high skilled men, which follows from that $P_{h,high} < P_{h,low}$ (see the implicit prices in Figure A.1). Since high skilled men face a higher wage rate in the market than do low skilled men, the lower implicit price for home production among high skilled men must be explained by the factor augmenting technologies. While low skilled men are estimated to be somewhat more productive at using capital, high skilled men are more efficient at using nondurable goods and their hours in home production (see Table 2.2). Clearly, the factors that bring down the implicit price of home production (z_n and z_l) dominate the factors that increase it (W and z_k).

That high skilled men consume more market services relative to nondurable goods than do low skilled men is thus explained by that they use fewer units of nondurable goods in relation to the quantity of home services. The ratio C_h/N is given by

$$\frac{C_{hjt}}{N_{jt}} = \left(\frac{P_{nt}}{P_{hjt}} \right)^{\sigma_{nq}} z_{njt}^{1-\sigma_{nq}} \tag{2.48}$$

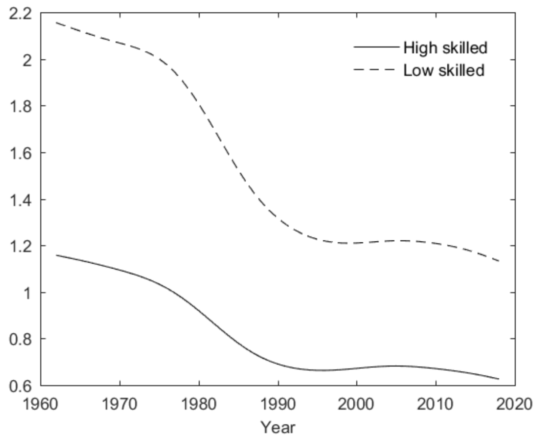

Figure A.1: *Implicit prices of home production: men*

Note: See the text for details about how these price indices are derived.

For the cross-sectional comparison, P_n drops out and we get

$$\frac{C_{h,high,t}/N_{high,t}}{C_{h,low,t}/N_{low,t}} = \left(\frac{P_{h,high,t}}{P_{h,low,t}}\right)^{-\sigma_{nq}} \left(\frac{z_{n,high,t}}{z_{n,low,t}}\right)^{1-\sigma_{nq}} \tag{2.49}$$

The ratio is higher among the high skilled for two reasons. First, their implicit price for home production is lower, as already mentioned. This increases the ratio, as it, ceteris paribus, implies that nondurable goods are relatively more expensive and therefore the household uses less of them.[21]

[21] Based on Equation (2.49), this might not be an obvious interpretation. Consider then instead Equation (2.48). There, a decrease in P_h, holding P_n constant, means that the implicit price of home services decreases relative to the price of nondurable goods. This does two things. First, it leads to the household using less nondurable goods, as it has become more expensive. Second, a decrease in the relative price of home services leads to more home services being produced. Both of these channels lead to increases in C_h/N.

Second, the ratio is affected by the nondurable goods augmenting technology, which is higher among the high skilled. Nondurable goods and the composite Q are gross complements in home production, and this makes the high skilled household substitute more towards the composite Q, thus decreasing N.

Let us look more closely at the drivers behind the non-monotonic development of C_m/N. For both types of households, the quantity ratio between market and home produced services decreased up until the early 1990s. Between then and 2018 it decreased, but the decrease was small. Looking back again at Equation (2.47), this is caused by a decrease and then a marginal increase of P_h. The decrease that lasted around three decades hence caused households to substitute away from market services and into home produced services. Wages did not decrease over this period and therefore the decrease must be driven by factor augmenting technologies and falling prices of capital and nondurable goods. With regards to technology, Figure 2.2 suggests that it is in particular the increases in z_n and z_k during this period that caused P_h to decrease.

An increase in C_h/N explains why C_m/N increased up until the 1990s. That is, in proportion to how much home services that were produced, the quantity of nondurable goods used decreased. This is driven both by that the price ratio P_n/P_h increased during this period and by the increase in z_n. From the early 1990s, the quantity of nondurable goods increased somewhat faster than the quantity of home services, leading the ratio C_h/N to fall, albeit modestly.

Up until the 1990s, households did substitute towards home services which, ceteris paribus, increased the demand for nondurable goods. But at the same time, households substituted away from nondurable goods in home production, to a large extent as a result of that their becoming more efficient at using it. These two forces counteracted each other, but the effect coming from substitution away from nondurable goods in home production dominated. From the beginning of the 1990s, the effects went in the same direction: households have been substituting somewhat back towards market services and have also become more reliant on nondurable goods in home production. These effects have caused the ratio C_m/N to decrease during the last three decades.

2.A.6 The implicit price of the consumption index C

Market services can be written as a function of home produced services:

$$C_m = P_h^{\sigma_{mh}} C_h \tag{2.50}$$

If I use this in the consumption index, I get

$$C = \left(C_m^{\frac{\sigma_{mh}-1}{\sigma_{mh}}} + C_h^{\frac{\sigma_{mh}-1}{\sigma_{mh}}} \right)^{\frac{\sigma_{mh}}{\sigma_{mh}-1}} \tag{2.51a}$$

$$C = \left(1 + P_h^{\sigma_{mh}-1} \right)^{\frac{\sigma_{mh}}{\sigma_{mh}-1}} C_h \tag{2.51b}$$

$$\implies C_h = \left(1 + P_h^{\sigma_{mh}-1} \right)^{\frac{\sigma_{mh}}{1-\sigma_{mh}}} C \tag{2.51c}$$

The total cost for one unit of C is

$$\text{cost}_C = C_m + P_h C_h \tag{2.52a}$$

$$= P_h^{\sigma_{mh}} C_h + P_h C_h \tag{2.52b}$$

$$= \left(1 + P_h^{\sigma_{mh}-1} \right) P_h C_h \tag{2.52c}$$

Substitute for C_h using the above expression of C_h as a function of C

$$\mathrm{cost}C = \left(1 + P_h^{\sigma_{mh-1}}\right) P_h \left(1 + P_h^{\sigma_{mh}-1}\right)^{\frac{\sigma_{mh}}{1-\sigma_{mh}}} C \tag{2.53a}$$

$$= \left(1 + P_h^{\sigma_{mh-1}}\right)^{\frac{1}{1-\sigma_{mh}}} P_h C \tag{2.53b}$$

$$= \left(1 + P_h^{1-\sigma_{mh}}\right)^{\frac{1}{1-\sigma_{mh}}} C \tag{2.53c}$$

So the implicit price index for the consumption index C is

$$P = \left(1 + P_h^{1-\sigma_{mh}}\right)^{\frac{1}{1-\sigma_{mh}}} \tag{2.54}$$

2.A.7 Explaining the Fall in Total Hours (and Why $\gamma > 1$ is Needed)

Total hours worked in the model are given by

$$L_{jt} = \left(\frac{W_{jt}}{P_{jt}}\right)^{\frac{1-\gamma}{\gamma+\varphi}} \chi_j^{-\frac{1}{\gamma+\varphi}} \tag{2.55}$$

where P is the implicit price of the consumption index C which is shown to be equal to (see Appendix 2.A.6)

$$P_{jt} = \left(1 + P_{hjt}^{1-\sigma_{mh}}\right)^{\frac{1}{1-\sigma_{mh}}} \tag{2.56}$$

where (see Appendix 2.A.5)

$$P_{hjt} = \left[\left(\frac{P_{nt}}{z_{njt}}\right)^{1-\sigma_{nq}} + P_{qjt}^{1-\sigma_{nq}}\right]^{\frac{1}{1-\sigma_{nq}}} \tag{2.57}$$

and (see Appendix 2.A.1)

$$P_{qjt} = \left[\left(\frac{P_{kt}}{z_{kjt}} \right)^{1-\sigma_{kl}} + \left(\frac{W_{jt}}{z_{ljt}} \right)^{1-\sigma_{kl}} \right]^{\frac{1}{1-\sigma_{kl}}} \tag{2.58}$$

I will show that the long-run trend in W/P is unambiguously increasing and that it is therefore necessary to have $\gamma > 1$ to match the decline in total hours worked. First, rewrite Equation (2.58):

$$P_{qjt} = \left[\left(\frac{P_{kt}}{z_{kjt} W_{jt}} \right)^{1-\sigma_{kl}} + \left(\frac{1}{z_{ljt}} \right)^{1-\sigma_{kl}} \right]^{\frac{1}{1-\sigma_{kl}}} W_{jt} \tag{2.59}$$

Now use Equation (2.59) in Equation (2.57) and shift the position of W:

$$\begin{aligned} P_{hjt} = \Bigg[\left(\frac{P_{nt}}{z_{njt}} \right)^{1-\sigma_{nq}} &+ \Bigg(\Bigg[\left(\frac{P_{kt}}{z_{kjt} W_{jt}} \right)^{1-\sigma_{kl}} \\ &+ \left(\frac{1}{z_{ljt}} \right)^{1-\sigma_{kl}} \Bigg]^{\frac{1}{1-\sigma_{kl}}} W_{jt} \Bigg)^{1-\sigma_{nq}} \Bigg]^{\frac{1}{1-\sigma_{nq}}} &\text{(2.60a)} \\ = \Bigg[\left(\frac{P_{nt}}{z_{njt} W_{jt}} \right)^{1-\sigma_{nq}} & \\ + \Bigg[\left(\frac{P_{kt}}{z_{kjt} W_{jt}} \right)^{1-\sigma_{kl}} &+ \left(\frac{1}{z_{ljt}} \right)^{1-\sigma_{kl}} \Bigg]^{\frac{1-\sigma_{nq}}{1-\sigma_{kl}}} \Bigg]^{\frac{1}{1-\sigma_{nq}}} W_{jt} &\text{(2.60b)} \\ = \Pi_{jt} W_{jt} & &\text{(2.60c)} \end{aligned}$$

where $\Pi_{jt} \equiv \left[\left(\frac{P_{nt}}{z_{njt} W_{jt}} \right)^{1-\sigma_{nq}} + \left[\left(\frac{P_{kt}}{z_{kjt} W_{jt}} \right)^{1-\sigma_{kl}} + \left(\frac{1}{z_{ljt}} \right)^{1-\sigma_{kl}} \right]^{\frac{1-\sigma_{nq}}{1-\sigma_{kl}}} \right]^{\frac{1}{1-\sigma_{nq}}}$.

At this point, note that the partial derivatives of Π with respect to

$\frac{P_n}{z_n W}$, $\frac{P_k}{z_k W}$ and $\frac{1}{z_l}$ are all positive. Since P_n and P_k fall over time while W, z_n, z_k and z_l increase (the latter three via the restriction), it is the case that Π falls over time. Now, insert it into the expression for P from Equation (2.56):

$$P_{jt} = \left(W_{jt}^{\sigma_{mh}-1} + \Pi_{jt}^{1-\sigma_{mh}}\right)^{\frac{1}{1-\sigma_{mh}}} W_{jt} \tag{2.61}$$

and thus:

$$\frac{W_{jt}}{P_{jt}} = \left(W_{jt}^{\sigma_{mh}-1} + \Pi_{jt}^{1-\sigma_{mh}}\right)^{\frac{1}{\sigma_{mh}-1}} \tag{2.62}$$

From this it is easy to see that $\frac{\partial \frac{W}{P}}{\partial W} > 0$ and $\frac{\partial \frac{W}{P}}{\partial \Pi} < 0$. Since W rises while Π falls over time, the ratio $\frac{W}{P}$ will increase. Hence from Equation (2.55) it is easy to see that it is necessary that $\gamma > 1$, i.e., that the income effect exceeds the substitution effect, for total hours to decrease.

Understanding how the relative prices P_n and P_k affect total hours is straightforward: a fall in these prices implies a fall in the prices of the inputs in home production, which necessarily leads to the implicit price of home production and the implicit price of the consumption index C to fall. A change in the wage rate does not only affect the numerator but also the denominator, as the household uses hours of domestic work, L_h, as an input in home production. That is, an increase in the wage rate raises the benefit of working in the market through the fact that the wage increases, but simultaneously it has a negative effect on the real wage since it raises the implicit price of home production. I can show that the first effect dominates the latter by, first, using Equation (2.62) to find that

$$\frac{d\frac{W}{P}}{W} = A_{jt}\left[W_{jt}^{\sigma_{mh-2}} - \Pi^{-\sigma_{mh}}\frac{d\Pi_{jt}}{W_{jt}}\right] \tag{2.63}$$

where $A_{jt} = \left(W_{jt}^{\sigma_{mh}-1} + \Pi_{jt}^{1-\sigma_{mh}}\right)^{\frac{1}{\sigma_{mh}-1}-1}$. From Equation (2.60b) it is easy to see that $\frac{d\Pi_{jt}}{dW_{jt}} < 0$, so therefore $\frac{d\frac{W}{P}}{W} > 0$.

2.A.8 The Implicit Price of the Composite Q for Women

The derivation of the implicit price P_q for women includes making a small but important adjustment compared to how it was derived for men. For women, the total cost of producing Q is:

$$\text{cost}_Q = P_k K + \frac{\kappa_h}{\kappa_m} W L_h \tag{2.64}$$

Note, importantly, that the ratio κ_h/κ_m multiplies the market wage rate. The total cost, and later on the implicit price, does not only capture the monetary cost of producing Q, but also captures a direct effect on utility. That it is κ_h/κ_m that multiplies the wage rate is simply due to the fact that it is κ_h/κ_m times as painful to work in home production compared to working in the market.

Substituting for L_h using Equation (2.21):

$$\text{cost}Q = \left[P_k + \frac{\kappa_h}{\kappa_m} W \left(\frac{P_k}{W}\frac{\kappa_m}{\kappa_h}\right)^{\sigma_{kl}} \left(\frac{z_k}{z_L}\right)^{1-\sigma_{kl}}\right] K \tag{2.65a}$$

$$= \left[1 + \left(\frac{P_k}{W}\frac{z_L}{z_k}\frac{\kappa_m}{\kappa_h}\right)^{\sigma_{kl}-1}\right] P_k K \tag{2.65b}$$

Now substitute for K using Equation (2.7) (which looks the same as for men) and then use the definition of Ω from Equation (2.23)

$$\mathrm{cost}_Q = \Omega P_k \Omega^{\frac{\sigma_{kl}}{1-\sigma_{kl}}} z_k^{-1} Q \tag{2.66a}$$

$$= \Omega^{\frac{1}{1-\sigma_{kl}}} \frac{P_k}{z_k} Q \tag{2.66b}$$

And hence the implicit price is

$$P_q = \Omega^{\frac{1}{1-\sigma_{kl}}} \frac{P_k}{z_k} \tag{2.67}$$

2.A.9 The Ratio K/N for Women

Start with the decomposition:

$$\frac{K_{jt}}{N_{jt}} = \frac{K_{jt}}{Q_{jt}} \times \frac{Q_{jt}}{N_{jt}} \tag{2.68}$$

The ratio K/Q is lower among high skilled women as compared to low skilled women, as a result of z_l being is higher for the prior type (see Appendix 2.A.2). Since K/N is higher among high skilled women than among low skilled women, Q/N must hence be higher for high skilled. What is it that causes this? This ratio is determined by

$$\frac{Q_{jt}}{N_{jt}} = \Omega_{jt}^{\frac{\sigma_{nq}}{\sigma_{kl}-1}} \left(\frac{P_{nt}}{P_{kt}}\right)^{\sigma_{nq}} z_{kjt}^{\sigma_{nq}} z_{njt}^{1-\sigma_{nq}} \tag{2.69}$$

Use that the implicit price of the composite Q for women that is derived in Appendix 2.A.8 is $P_q = \Omega^{\frac{1}{1-\sigma_{kl}}} \frac{P_k}{z_k}$, which yields

$$\frac{Q_{jt}}{N_{jt}} = \left(\frac{P_{nt}}{P_{qjt}}\right)^{\sigma_{nq}} z_{njt}^{1-\sigma_{nq}} \tag{2.70}$$

So the cross-sectional difference is given by

$$\frac{Q_{high,t}/N_{high,t}}{Q_{low,t}/N_{low,t}} = \left(\frac{P_{q,high,t}}{P_{q,low,t}}\right)^{-\sigma_{nq}} \left(\frac{z_{n,high,t}}{z_{n,low,t}}\right)^{1-\sigma_{nq}} \tag{2.71}$$

On average over all the years, the ratio Q/N has been 2.45 times as large for high skilled women as for low skilled women. Although the implicit price of the composite Q is significantly lower for high skilled women as compared to low skilled women and has a positive effect on the ratio under study, quantitatively the effect is negligible, as a result of the low value of σ_{nq}. The difference that explains the absolute majority of why the ratio Q/N is significantly higher among high skilled women is that they are more efficient at using nondurable goods when producing home services. As the composite Q and nondurable goods are gross-complements, this makes the household substitute towards using more Q, thus increasing the demand for capital. Hence, it can be summarized as the main reason why the ratio K/N being higher among high skilled women than among low skilled women is that high skilled women use nondurable goods more efficiently and, since the factors of production are gross complements, substitute towards the composite Q, which is partly made up of capital.

The increase in the capital-to-nondurable goods ratio between the years 1962 and 2018 is explained by both ratios K/Q and Q/N increasing. But like for men, Q/N was the dominant force behind the increase up until the early 1990s, from which point in time K/Q took over as the driving force. Up until the 1990s, labor augmenting technology grew at a rapid pace, while the relative price P_n/P_q decreased. While the first of these developments contributed positively to the increase in Q/N, the latter had the opposite effect. Primarily as a result of the low value of σ_{nq}, the increase in z_n dominated. The growth in

z_n stopped in the middle of the 1990s and the ratio Q/N decreased slightly.

The modest increase in K/Q up until the 1990s was followed by a more pronounced increase thereafter, fueling the increase in the capital-to-nondurable goods ratio. Like for men in Appendix 2.A.4, the ratio K/Q can be written as

$$\frac{K_{jt}}{Q_{jt}} = \left(\frac{P_{jt}}{P_{kt}/z_{kjt}}\right)^{\sigma_{kl}} z_{kjt}^{-1} \tag{2.72}$$

The first part on the right-hand-side of Equation (2.72), $\frac{P_{qjt}}{P_{kt}/z_{kjt}}$, grew monotonically over time, but the speed at which it grew was higher from the 1990s onwards. Moreover, the lower growth rate of this ratio before the 1990s was also counteracted by the growth in z_k that comes in as the last term in Equation (2.72) with a negative effect. I will again, like in Appendix (2.A.4), use that

$$\frac{P_{qjt}}{P_{kt}/z_{kjt}} = \Omega_{jt}^{\frac{1}{1-\sigma_{kl}}} \tag{2.73}$$

where

$$\Omega_{jt}^{\frac{1}{1-\sigma_{kl}}} = \left[1 + \left(\frac{z_{kjt}W_{jt}}{z_{ljt}P_{kt}}\frac{\kappa_{ht}}{\kappa_{mt}}\right)^{1-\sigma_{kl}}\right]^{\frac{1}{1-\sigma_{kl}}} \tag{2.74}$$

There are three ratios inside the expression that change over time: z_k/z_L, W/P_k, and κ_h/κ_m. From equations (2.72) and (2.74) it is easy to see that K/Q increases with all of these ratios.

I produce three counterfactuals to gauge the importance of the developments over time in (i) technology, (ii) the relative price W/P_k, and (iii) the ratio κ_h/κ_m, respectively. In each counterfactual, I fix them, one at the time, at their 1962 levels. The results are illustrated in Figure A.2. From both panels, it is obvious that the increases would have been significantly more positive had z_k/z_l stayed at its initial level. Since we know that the ratio z_k/z_l increased over time and had a positive effect on K/Q, this effect from shutting down developments in z_k and z_l must come from the term z_k^{-1} in Equation (2.72). Absent the positive developments in either W/P_k or κ_h/κ_m, the K/Q ratio would barely have changed between the years 1962 and 2018. The contribution from W/P_k has been somewhat larger than that from κ_h/κ_m after the year 2000, but the more striking result is that, since the year 1962, the contributions from the two were similar. In summary: increases in both κ_h/κ_m and W/P_k explain, by roughly equal proportions, the increases in K/Q since the year 1962.

2.A.10 The Implicit Price of Home Services and C_m/N for Women

With the implicit price of the composite Q having already been derived for women in Appendix 2.A.8, the remaining steps for deriving the implicit price of home services for women are identical to those for men. Moreover, the final expression for the implicit price of home services will look like it does for men (see Appendix 2.A.6):

$$P_{hjt} = \Gamma_{jt}^{\frac{1}{1-\sigma_{nq}}} \frac{P_{nt}}{z_{njt}} \tag{2.75}$$

only now Γ_{jt} contains P_{qjt}, in which κ_h/κ_m enters. For convenience, I write down the decomposition here again:

$$\frac{C_{mjt}}{N_{jt}} = \frac{C_{mjt}}{C_{hjt}} \times \frac{C_{hjt}}{N_{jt}} \tag{2.76}$$

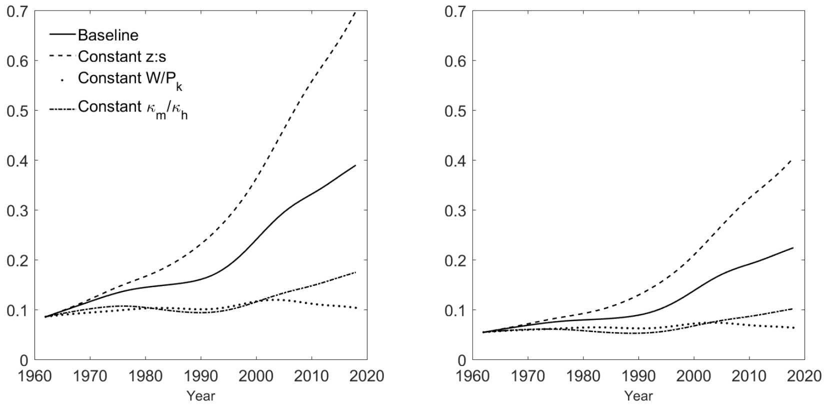

(a) Low skilled women *(b) High skilled women*

Figure A.2: *The* K/Q *ratio for single women in some alternative scenarios*

Note: *Panel (a)*: The solid line represents the baseline development in the model. In each of the three alternative scenarios, some of the prices or parameters are kept constant at their 1962 levels. The dashed line represents the development of the ratio when z_k and z_L are held constant; the dotted line represents the development when W/P_k is held constant; the dashed-dotten line represents the development for when κ_m/κ_h is held constant. *Panel (b)*: Same as in panel (a), but for high-skilled women.

To understand why the ratio under study is higher among high skilled women, I follow the same steps as I did when analyzing men. The overall conclusions about why high skilled women have a higher ratio than do low skilled women are the same as for men. The first ratio in the decomposition from Equation (2.76) is larger for low skilled women, as a result of them producing significantly less home services compared to high skilled women. High skilled women also use greater quantities of non-durable goods in home production than do low skilled women, but still the ratio C_h/N is higher among high skilled women, which explains why the ratio C_m/N is higher in this group.

The ratio between market services and home produced services is given by

$$\frac{C_{mjt}}{C_{hjt}} = P_{hjt}^{\sigma_{mh}} \tag{2.77}$$

Through the lens of the model, C_m/C_h is higher among low skilled women because of a higher implicit price of home services (see the implicit prices in Figure A.3). As was shown in Appendix 2.A.5, the implicit price increases with W and decreases with z_i $\forall i \in \{n, k, l\}$. Since high skilled women face higher wage rates than do low skilled women, the higher implicit price of home services for low skilled must be driven by differences in factor augmenting technologies. Like for men, the efficiency at which high skilled women use capital is lower than among low skilled women, but high skilled women are more efficient in their usage of nondurable goods and hours of domestic work. That low skilled use nondurable goods and hours in home production less efficiently explains why their implicit price of home services is higher and why they consume more market services relative to home services.

Like for men, the ratio C_h/N is given by

$$\frac{C_{jt}}{N_{jt}} = \left(\frac{P_{nt}}{P_{hjt}}\right)^{\sigma_{nq}} z_{njt}^{1-\sigma_{nq}} \tag{2.78}$$

and the between-household differences by

$$\frac{C_{h,high,t}/N_{high,t}}{C_{h,low,t}/N_{low,t}} = \left(\frac{P_{h,high,t}}{P_{h,low,t}}\right)^{-\sigma_{nq}} \left(\frac{z_{n,high,t}}{z_{n,low,t}}\right)^{1-\sigma_{nq}} \tag{2.79}$$

The ratio on the left-hand-side has been 2.43 on average. Even though

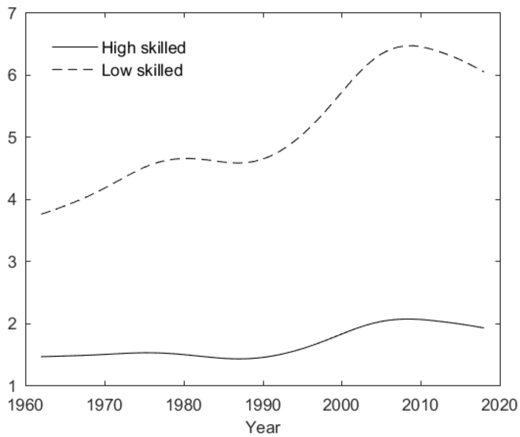

Figure A.3: *Implicit prices of home production: women*

Note: See the text for details about how these price indices are derived.

the differences in the implicit prices of home services are quantitatively significant: the average of $P_{h,high}/P_{h,low}$ over all the years has been 0.33, the low value of σ_{nq} leads to a small effect on the ratio on the left-hand-side. It is the fact that high skilled women use nondurable goods more productively that explains why the ratio C_h/N is higher among them, since households substitute away from nondurable goods and towards the composite Q when z_n increases. As can be seen in Figure 2.6c, the increase in the ratio C_m/N over time is severely over-predicted by the model. Moreover, a decline in the ratio for low skilled women started at the beginning of the 1990s while the model predicts a commenced, modest, decline more than a decade later. Nevertheless I will here undertake an analysis of the initial increase and then subsequent decrease in the market services to nondurable goods ratio.

The story behind the increase in the ratio C_m/N over time is qualitatively different compared to the story for men. For men, the increase was generated by substitution away from nondurable goods in home production, but the increase was quantitatively dampened by substitution away from market services and to home services. For women, the increase in the ratio that is generated by the model up until the first years of the 1990s is explained by *both ratios* on the right-hand-side of Equation (2.76) increasing. After that point in time, the increase that is generated by the model, lasting until around the year 2010, is driven by continued substitution away from home produced services and towards market produced services, i.e., an increase in C_m/C_h. The other ratio, C_h/N, starts decreasing during the early/mid 1990s, mechanically explained by the fact that the raw quantity of nondurable goods that households use in home production decreases relative to the quantity of home services produced.

Continuous increases the implicit price of home services up until around the year 2010 caused households to substitute from home services to market services. Knowing that these implicit prices decreased for men, what lead to them increase for women? From Appendix 2.A.5:

$$P_{hjt} = \left[\left(\frac{P_{nt}}{z_{njt}} \right)^{1-\sigma_{nq}} + P_{qjt}^{1-\sigma_{nq}} \right]^{\frac{1}{1-\sigma_{nq}}} \tag{2.80}$$

The implicit price increases with both the price per efficiency units of nondurable goods, P_n/z_n, and the implicit price of the composite Q, P_q. The first of these two fell up until the early 2000s and from there onwards only changed marginally. Hence, increases in the implicit prices of home services were driven by increases in P_q. Building on Equation (2.67) from Appendix 2.A.8, the implicit price of the composite Q can be written as

$$P_{qjt} = \left[\left(\frac{P_{kt}}{z_{kjt}} \right)^{1-\sigma_{kl}} + \left(\frac{W_{jt}}{z_{ljt}} \frac{\kappa_{ht}}{\kappa_{mt}} \right)^{1-\sigma_{kl}} \right]^{\frac{1}{1-\sigma_{kl}}} \tag{2.81}$$

The first term inside the brackets, P_k/z_k, decreases monotonically with time. The price per efficiency units of hours of domestic work, W/z_l first increases first up until around the mid 1970s and then falls back. The development is, however, very similar to the development for men, as the development of z_l in its indexed form is identical and the changes in wage rates are similar.

Changes in κ_h/κ_m over time are responsible for the increase in P_q. Figure A.4 compares the developments of the implicit prices for low and high skilled women, respectively, in the model with baseline parameters and also illustrates the counterfactual developments, had the ratio κ_h/κ_m stayed constant at its 1962 level. Clearly, the increases in implicit prices would not have occurred without κ_h/κ_m increasing over time.

A couple of years before 2010, the ratio κ_h/κ_m stopped increasing. The ratio W/z_l fell for both low skilled and high skilled women, and P_k/z_k continued falling. With all of these components either decreasing or staying more-or-less constant, the implicit price indices for both types of women decreased, leading to some substitution back into home produced services.

Let me summarize the development of C_m/C_h and the drivers behind it. As the implicit prices of home services for both types of women increased up until 2009, they substituted towards market services and away from home services. After 2009, these implicit prices fell, leading to substitution in the other direction. The increases in im-

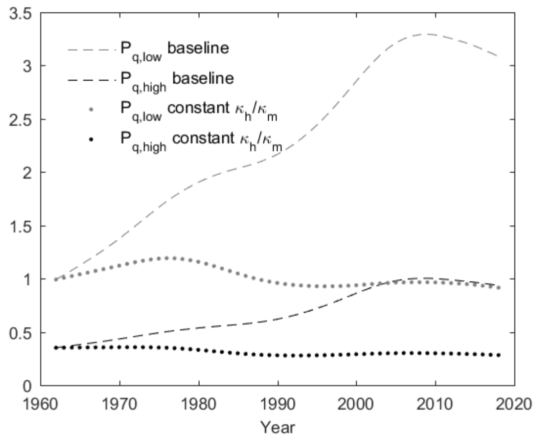

Figure A.4: *Counterfactual development of* P_q *for women*

Note: *Dashed lines*: developments of implicit prices of the composite Q for low skilled women (gray) and high skilled women (black) in the model with baseline parameters. *Dotted lines*: counterfactual development when holding κ_h/κ_m fixed at its 1962 level. All lines are indexed to the value of the baseline implicit price for high skilled women in the year 2010.

plicit prices that lasted for many decades were driven by an increase in κ_h/κ_m, i.e., that the distaste for working at home increased relative to the distaste for working in the market. After the year 2009, this ratio stopped increasing. As primarily P_k/z_k continued to decrease also after 2009, the implicit prices fell.

Regarding the changes over time in C_h/N, the developments are similar to the developments for men. The ratio increased until 1992 and then fell back. In the model, the ratio is determined by

$$\frac{C_{hjt}}{N_{jt}} = \left(\frac{P_{nt}}{P_{hjt}}\right)^{\sigma_{nq}} z_{njt}^{1-\sigma_{nq}} \tag{2.82}$$

Although the decrease in $\frac{P_n}{P_h}$ was larger than the increase in z_n until 1992, the low value for σ_{nq} leads to the latter effect still dominating, in turn giving the increase in C_h/N over this period. After 1992, z_n stopped increasing, even decreasing somewhat, while P_n/P_h continued to decrease. With two factors joining forces in the negative direction, they both caused C_h/N to decline. In summary: C_h/N increased up until the year 1992 as a result of nondurable goods being used more efficiently over time, causing households to substitute away from it and into the composite Q, since Q and N are gross complements. This increase in z_n stopped after the year 1992, while the relative price P_n/P_h continued to decrease. The latter decrease had an effect on the ratio in question as nondurable goods became more affordable as an input in home production, but still, because of the low value of σ_{nq}, only had a small quantitative effect. I do not show it, but a quick investigation into the effect of κ_h/κ_m on this ratio revealed that it was not so important, again because of the low value of σ_{nq}.

2.A.11 Time Allocation Among Women

For women, the equation that I use in the analysis looks as follows:

$$\kappa_{mt}L_{hjt} + \kappa_{ht}L_{hjt} = \left(\frac{W_{jt}}{P_{jt}\kappa_{mt}}\right)^{\frac{1-\gamma}{\gamma+\varphi}} \chi_j^{-\frac{1}{\gamma+\varphi}} \tag{2.83}$$

where P is the implicit price index for the consumption index C (see Appendix 2.A.6). Clearly, due to the existence of κ_m and κ_h, the analysis is more difficult compared to the analysis for men. A direct and convenient method for understanding the cross-sectional differences and developments over time is to vary prices and parameters and look at counterfactual outcomes.

I start by examining the role played by the wage rate W. Note, importantly, that the price index P depends on W, which furthers complicates the analysis. I construct two counterfactual scenarios: one in which I shift the levels of the wage paths, and a second in which I keep wages fixed at their 1962 levels.

In the first scenario, I increase the wage rates for low skilled in all years by 63 percent, which corresponds to the level difference in wage rates between high and low skilled in the year 1962. I make the same type of adjustment to wages among high skilled, shifting them down, such that the wage rate in 1962 is the same as for low skilled women. The results are shown in Figure A.5. Interestingly, the differences in time-use would have been larger with a smaller wage-gap between high and low skilled women; the exercise shows that (i) a higher wage rate depresses market hours, and (ii) increases hours in home production.

To grasp the importance of wages for the development over time, the wage rates are held at the 1962 levels in the second scenario. See Figure A.6 for results. For high skilled women, the development of the wage rate was essentially unimportant until the beginning of the 1990s, which is captured by the fact that the solid and dotted lines lie on top of each other until that point in time. Over the next coming 15 years, up until around the year 2005, the increase in the wage rate led to falling market hours, which is inferred from the observations that market hours would have continued to rise over this period with constant wage rates. For low skilled, it is clear that the increase in the wage rate over time depressed market hours and also had a negative net effect on total hours. Market hours decrease by around -20% in the model, but the decrease would have been significantly smaller

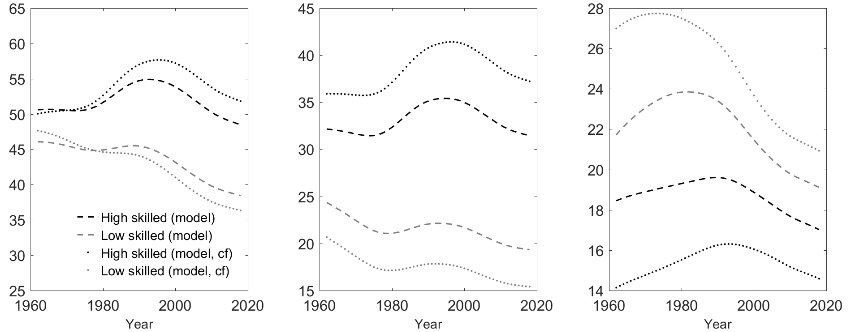

(a) *Total* ***(b)*** *Market* ***(c)*** *Home*

Figure A.5: *time-use among women, shifting wage levels*

Note: Measured in units of hours per week. *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production.

absent the increase in the wage rate.

How would things have changed if social norms, i.e., κ_m and κ_h, had stayed constant from the year 1962? The model's prediction for this is shown in Figure A.7. Had social norms not changed, women would have increased the hours in home production, which is clearly in contrast to the decrease that is observed in the data and reproduced by the model with baseline parameters. In the absence of changes in social norms, the market hours among both low and high skilled women would have decreased. The effect is large for high skilled, for whom market hours would have fallen by four hours between the years 1962 and 2018. Changes in social norms were clearly important for why market hours increased in most years between 1980 and 2000. More specifically, the underlying cause for the increase during these years is driven by a substantial decline in the relative distaste for

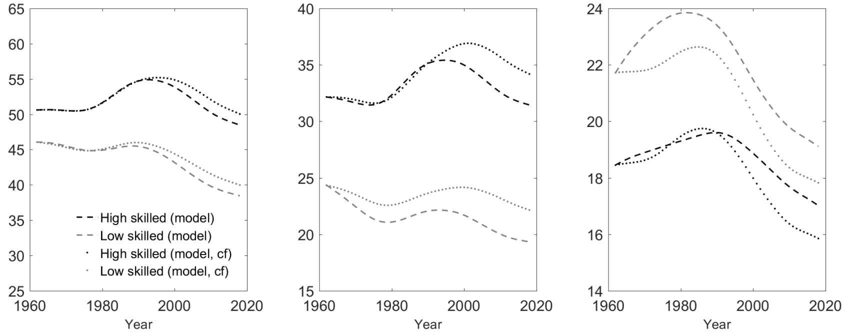

(a) Total *(b) Market* *(c) Home*

Figure A.6: *time-use among women, with constant wages*

Note: Measured in units of hours per week. *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production.

market work (see Figure 2.5b).

I adjust factor augmenting technologies in the last two counterfactual exercises. First, I hold each respective factor augmenting technology parameter, for each type of woman, fixed at their levels from 1962 (see Figure A.8). Some differences between the baseline model output and output from the counterfactual are clearly visible. However, all-in-all, shutting down this type of growth has limited effects on time-use. Second, I adjust each respective factor augmenting technology parameter for low skilled women to the level estimated for high skilled, and vice versa. As can be seen in Figure A.9, this has considerable effects on both levels and trends. Market hours increase substantially for low skilled, and decrease by similar magnitudes for high skilled. Low skilled women decrease the hours that they work at home, while they increase among high skilled. The level change in

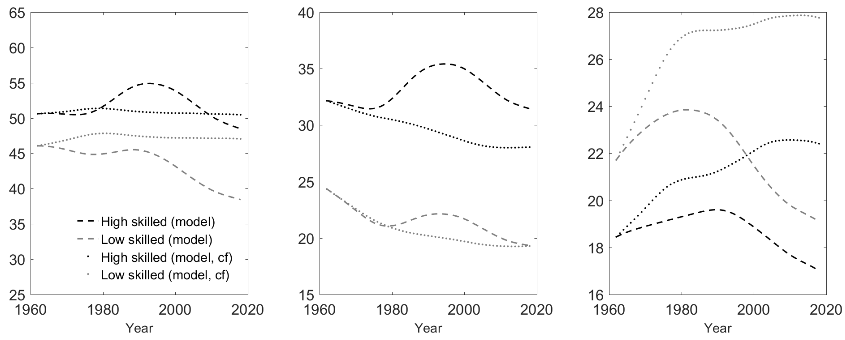

(a) *Total* ***(b)*** *Market* ***(c)*** *Home*

Figure A.7: *time-use among women, with constant social norms*

Note: Measured in units of hours per week. *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production.

productivites also affects the trends over time. Underlying this result is the non-linear nature of a nested CES structure, like the one employed in the current paper. Clearly, the changes over time in each respective factor augmenting technology, κ_h and κ_m, and the wage rates interact with the levels of the factor augmenting technology parameters in a way that has important implications for the trends.

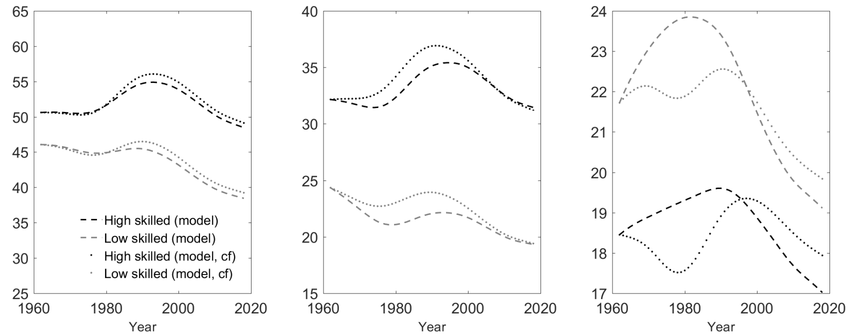

(a) Total *(b) Market* *(c) Home*

Figure A.8: *Time-use among women, without productivity growth*

Note: Measured in units of hours per week. *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production.

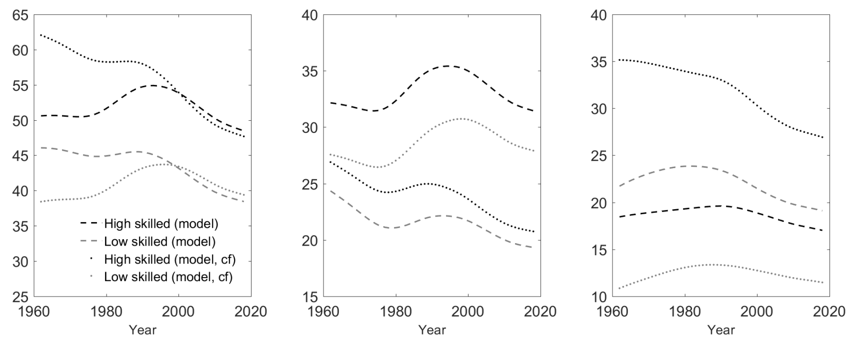

(a) Total *(b) Market* *(c) Home*

Figure A.9: *Time-use among women, shifting productivity levels*

Note: Measured in units of hours per week. *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production.

2.B Figure Appendix

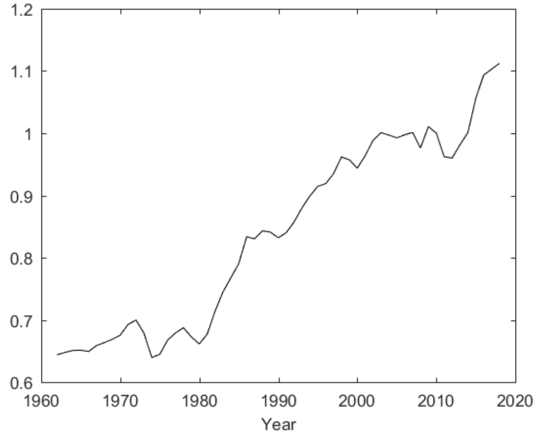

Figure B.1: *Relative price of market services: unfiltered*

Note: Relative price between market services and nondurable goods.

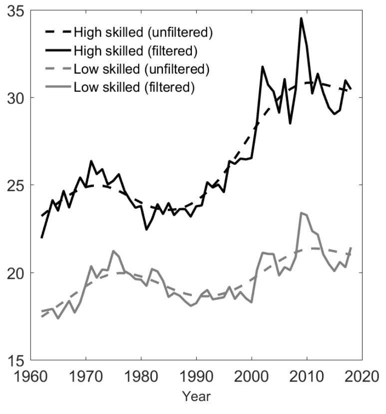

(a) Men

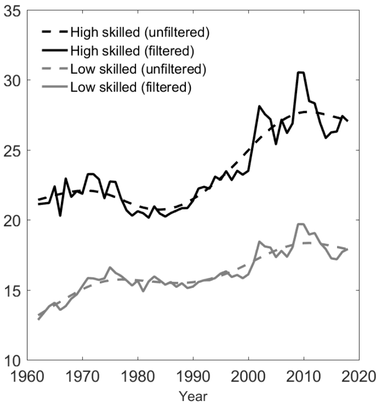

(b) Women

Figure B.2: *Wage rates*

Note: *Panel (a)*: Development of wage rates for low and high skilled men. *Panel (b)*: Development of wage rates for low and high skilled women. Solid lines indicate unfiltered data, while the data series in dashed are filtered using the HP-filter.

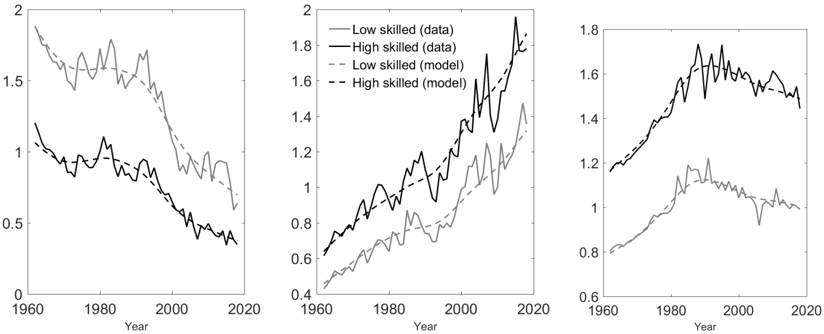

(a) $\mathsf{L_h/K}$ ***(b)*** $\mathsf{K/N}$ ***(c)*** $\mathsf{C_m/N}$

Figure B.3: *Quantity ratios for men: model and unfiltered data*

Note: *Panel (a)*: Ratio between hours of domestic work and capital. *Panel (b)*: Ratio between capital and nondurable goods. *Panel (c)*: Ratio between market services and nondurable goods. The series in each respective panel are normalized to one for low skilled men in the year 2010.

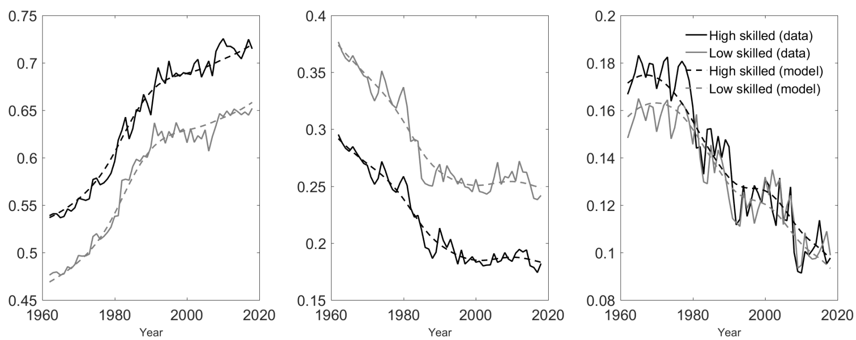

(a) *Market services* ***(b)*** *Nondurable goods* ***(c)*** *Capital*

Figure B.4: *Expenditure shares for men: model and unfiltered data*

Note: *Panel (a)*: Expenditure shares on market services *Panel (b)*: Expenditure shares on nondurable goods *Panel (c)*: Expenditure shares on capital/durable goods.

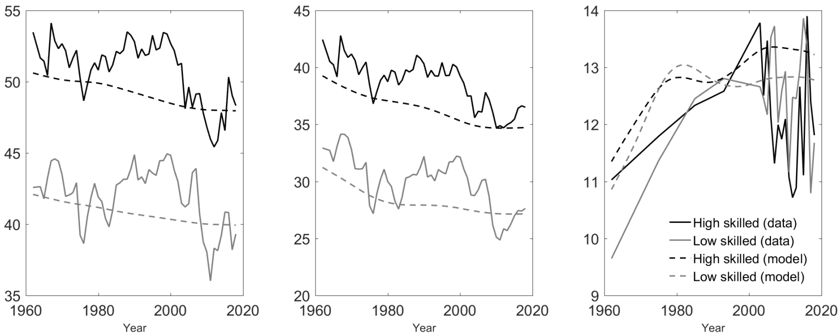

(a) *Total* ***(b)*** *Market* ***(c)*** *Home*

Figure B.5: *Hours worked for men: model and unfiltered data*

Note: *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production.

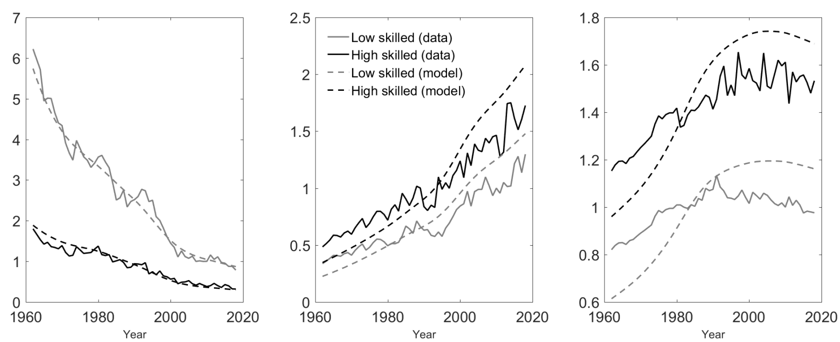

(a) $\mathsf{L_h/K}$ ***(b)*** $\mathsf{K/N}$ ***(c)*** $\mathsf{C_m/N}$

Figure B.6: *Quantity ratios for women: model and unfiltered data*

Note: *Panel (a)*: Ratio between hours in home production and capital. *Panel (b)*: Ratio between capital and nondurable goods. *Panel (c)*: Ratio between market services and nondurable goods. The series in each respective panel are normalized to one for low skilled women in year 2010.

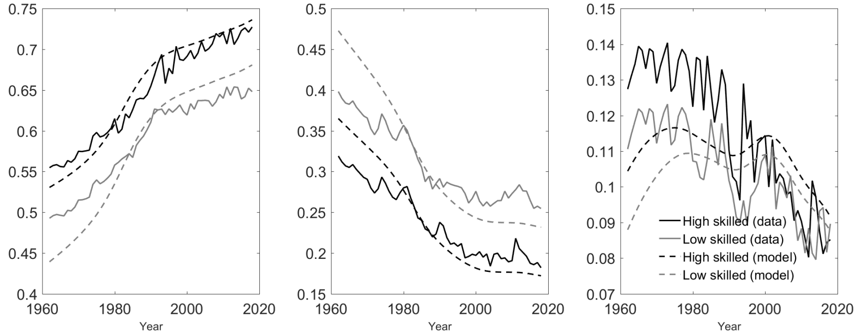

(a) *Market services* ***(b)*** *Nondurable goods* ***(c)*** *Capital*

Figure B.7: *Expenditure shares for women: model and unfiltered data*

Note: *Panel (a)*: Expenditure shares on market services *Panel (b)*: Expenditure shares on nondurable goods *Panel (c)*: Expenditure shares on capital/durable goods.

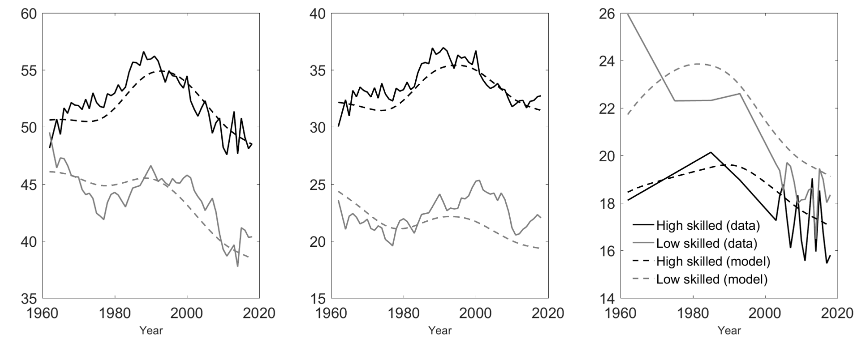

(a) *Total* ***(b)*** *Market* ***(c)*** *Home*

Figure B.8: *Hours worked for women: model and unfiltered data*

Note: *Panel (a)*: Total work hours (home plus market). *Panel (b)*: Hours worked in the market. *Panel (c)*: Hours worked in home production.

Chapter 3

It Runs in the Family: Occupational Choice and the Allocation of Talent*

*This paper was written jointly with John Kramer and Jósef Sigurdsson. We thank Timo Boppart, Konrad Burchardi, Gustavo de Souza, Mitch Downey, Per Krusell, Kieran Larkin, Kurt Mitman, Jay Lee, Arash Nekoei, Peter Nilsson, David Seim, Fabian Sinn, Xueping Sun, Horng Chern Wong, and seminar participants at various institutions helpful feedback and suggestions. We acknowledge financial support from Jan Wallander and Tom Hedelius Foundation.

3.1 Introduction

When children choose which occupation to pursue when they grow up, they often follow in their parents' footsteps. Sons of medical doctors, for example, are disproportionately more likely than others to become doctors themselves (Becker, 1959; Lentz and Laband, 1989). The same holds true for a range of occupations, from farmers to fishermen, bankers to businessmen. This tendency for intergenerational continuity of occupation—or occupational inheritance—has long drawn the attention of social scientists (Rogoff, 1953; Blau and Duncan, 1967). The fundamental, yet still unanswered, question is what the consequences of this phenomenon are for social mobility and the allocation of talent.

The implications of occupational persistence depend on its underlying drivers. If parents select into occupations based on abilities that are passed down to their children, persistence in occupations and inequality in incomes stem from differences in talents. If, by contrast, sorting is driven by exposure to, and information about, occupations, then intergenerational persistence in occupations and incomes may reflect lost opportunities and misallocation of talent.

In this paper, we study why jobs run in families and what the economic consequences of this are. More specifically, we investigate whether children choose to enter the same occupations as their parents due to sorting on common skills, or if parents affect their children's choices through other means, possibly distorting the efficient allocation of talents in the economy. Harnessing unique data on worker skills and personality traits, we use a machine-learning algorithm to measure occupational skill requirements and to quantify how well workers match to occupations based on their abilities.[1] Utilizing these results, we estimate a structural general equilibrium Roy (1951) model of costly occupational choice, which incorporates the notion that it may be less costly for children to choose their parent's occupation. Our central finding is that once this influence of parental background is removed and all individuals have equal opportunities, the propensity of a son to follow in the career footsteps of his father drops by half. This implies substantial misallocation in the baseline economy. However, despite this large reduction in occupational following, the impact on intergenerational mobility and aggregate productivity is small. Although workers relocate to occupations in which they have a comparative advantage, earnings and skill demands in the new occupations are similar to those of their fathers' occupations, implying only small earnings gains from resorting. Then, exploiting cross-sectional variation in the decline of fathers' occupations, we document that, on average, sons gain from not following into their father occupations, through higher earnings.

We begin by studying the occupational choices of children, documenting how they relate to the occupations of their parents and their skills. Using administrative data on the Swedish population, we show that children are disproportionately more likely to choose the same occupation as their parents, compared to the hypothetical alternative of random choices. This holds true even in narrowly defined occupations and across the whole spectrum of occupational categories.[2] In

[1] Unfortunately, since our data on skills originate in tests done in association with the Swedish military draft, which women were not subject to, most of our analysis is limited to the study of men. However, as we document, women have, if anything, a greater tendency to follow in the footstep of their parents, primarily their mothers.

[2] Our main sample is based on a classification of 91 occupations that is consistent from 1960 until today.

addition, we find that children who do not follow their parents into the same occupation often stay close to it, i.e., within the same broad occupational classification. Importantly, we document that occupational inheritance is a key driver of intergenerational persistence in earnings, while within-occupation earnings differences of individuals appear to be much less important. Our findings mirror those in an extensive literature within social sciences documenting occupational inheritance (e.g., Rogoff, 1953; Blau and Duncan, 1967; Laband and Lentz, 1985). We extend this work by quantifying the macroeconomic implications of occupational inheritance for intergenerational mobility and the allocation of talent in the economy.

A crucial ingredient for our empirical analysis and model calibration are estimates of each individual's potential earnings across all occupations. To obtain such estimates, we train a machine-learning algorithm to construct skill-based earnings predictions for all individuals in every occupation, utilizing detailed data on individuals' cognitive abilities and personality traits. The key idea underlying this approach is that each skill—or, more generally, each combination of skills—is differently productive in different occupations.[3] For each individual, our skill data consist of a vector of productive talents—both cognitive skills and personality traits–which was collected during tests

[3]The idea underlying our approach is similar to that of Gibbons and Waldman (2004) and Gathmann and Schönberg (2010), who assume that workers are endowed with a productivity in each task and then choose an occupation based on the tasks it requires. As first emphasized by Roy (1951), later by Sattinger (1975), and empirically documented in been documented in, e.g., Autor and Handel (2013) and Fredriksson et al. (2018), this leads workers to sort into occupations based on comparative advantage. This implies that the skills of incumbent workers can be used to measure the skill requirements of each occupation. Motivated by this, our approach is to train an algorithm on the earnings of incumbents in each occupation, excluding occupational followers, and predict earnings for every potential entrant.

and evaluations done in association with the Swedish military draft.[4] During our sample period, almost all men were subject to these tests between the ages of 18 and 19. Importantly, this implies that the traits are measured before individuals enter the labor market and can therefore not be directly influenced by the occupations they choose.

To study the implications of occupational following for the allocation of talent in the economy and for intergenerational income mobility, we propose a general equilibrium model of occupational choice. Into an otherwise standard Roy (1951) model, where individuals are endowed with a range of talents and occupations differ in their demands for these talents, we introduce utility costs of entering occupations and, moreover, explicitly account for the influence of the fathers' occupations on their sons' occupational choices. Having a father in an occupation acts as an advantage—or an entry cost "discount"—when entering the same occupation. In addition, we incorporate similar reductions in entry costs for individuals who enter the same broad occupational group as their fathers, accounting for the general tendencies we measure in the data. These cost discounts make choosing a father's occupation more attractive, or, equivalently, less costly, compared to individuals with the same abilities but from a different background. However, cost reductions will only influence occupational choice if sons are not already sorting into their fathers' occupations based on common comparative advantage. Our main experiment is the removal of these cost reductions, which constitutes a large deviation from the baseline economy. As a result, a model solved in general equilibrium, in which prices can adjust in response

[4]The cognitive traits are: inductive, verbal, spatial, and technical ability; the non-cognitive traits are: social maturity, intensity, psychological energy, and emotional stability. The non-cognitive traits are evaluated by trained psychologists (Mood et al., 2012)

to substantial changes in the occupational composition, is important for evaluating the aggregate effects.

Using the earnings predictions as an input for our model, we calibrate the entry costs, and their respective reductions, depending on a son's background, to match the patterns of occupational choice we measure in the data. Our model is able to generate the strong intergenerational persistence of occupations and, in addition, intergenerational persistence in earnings very similar to what we observe in the data. We find that the model-implied entry costs are highly correlated with factors such as educational and work experience requirements which vary across occupations. This is in line with occupational entry costs representing required investments in human capital prior to entry. The estimated cost discounts by paternal background may reflect a wide range of factors, including unequal access to information and education (e.g. medical doctors and lawyers (Lentz and Laband, 1989)), experience and knowledge (e.g. farmers, lawyers, and entrepreneurs (Laband and Lentz, 1983; Lentz and Laband, 1990; Laband and Lentz, 1992)), network and connections (e.g. corporate managers and politicians (Dal Bó et al., 2009)), and family firms (e.g. pharmacists (Mocetti, 2016)).[5] Finding suitable quasi-experimental variation in all factors behind occupational following is challenging, if at all possible. The attractiveness of our approach is that the structural model we present is able to account for the general phenomenon

[5]Prior literature has attempted to explicitly model channels through which parents affect the occupational choices of their children. Following the seminal work of Becker and Tomes (1986), Hellerstein and Morrill (2011) develop a model in which fathers vest their children with human capital that is specific to their own occupations. To the extent such investment is common across siblings, our results show that such investment cannot by themselves account for the patterns in the data. Lo Bello and Morchio (2021) model occupational persistence as reflecting transmission of skills, contacts and preferences.

of occupational following observed in the data, across all occupations, without requiring us to explicitly model all the possible factors that could give rise to it.

We use the model to conduct a counterfactual, in which we study the consequences of leveling the playing field of occupational entry. That is, we evaluate the impact of removing the heterogeneous entry-costs reductions based on fathers' occupations. This provides all workers who share the same skills with the same opportunities, irrespective of their parental background. The results are striking: close to 10 percent of individuals switch occupations; the overall tendency for sons to follow into the occupation of their father falls from 8.6 to 3.5 percent. Occupational following decreases in almost all occupations.

At first sight, our results imply that patterns of occupational choice reflect substantial misallocation of talent across occupations, as a large share of workers move to different occupations once occupational choice only depends on differences in abilities but not on background. Despite a substantial reallocation of workers across occupations, however, we measure only a modest increase in intergenerational earnings mobility, while aggregate income is unaffected. The intergenerational correlation in income ranks of fathers and sons decreases from about 0.291 to 0.245. The increase in mobility is most prominent in the bottom quintile of the father's income distribution. This reduction roughly corresponds to a fifth of the difference in income mobility between the US (Chetty et al., 2014) and Canada (Corak and Heisz, 1999) or Denmark (Boserup et al., 2013).

The reason for this small effect on income mobility is that although a large share of previous followers now choose different occupations than their fathers, they on average, depending on their father's income rank, select into occupations with similar incomes. For example, sons of doctors become finance and sales associates, and sons of lawyers become engineers. This reflects the fact that although, in the counterfactual economy, sons move to occupations that may utilize different parts of their skillsets, by choosing to follow their father, they were still relatively close to sorting into occupations in which they had a comparative advantage.

Next, we estimate the effect of following on sons' earnings. Using a difference-in-differences approach that exploits cross-cohort variation in employment growth in father's occupation, driven by, e.g., structural factors such as automation, we find that sons whose fathers' occupations are in decline are less likely to follow into that occupation. Moreover, as a consequence of choosing another occupation, the sons that *do not* follow their fathers due to such decline *gain* in terms of prime-age income. This is driven by sons with skills that are the least aligned with their fathers occupation, implying that by following they do not exploit their full economic potential. Our structural model yields similar conclusions: for individuals with fathers at the bottom of their income distribution, following leads to a stark decline in earnings. Individuals who follow a father whose earnings put him in the top decile of his distribution gain from following, compared to their other occupational options.

Our results build on, contribute to, and extend, several literatures. First, our results contribute to an extensive literature in economics and sociology documenting interegenerational persistence in occupations (e.g., Rogoff, 1953; Blau and Duncan, 1967; Laband and Lentz, 1985; Long and Ferrie, 2013), the voluminous empirical literature in economics measuring within-family income correlations across generations (see, e.g., Black et al., 2011, for survey), and the vast literature in sociology that has measured intergenerational mobility across occupational economic status levels (see, e.g., Ganzeboom et al., 1991, for survey). The work that has looked beyond intergenerational correlations has primarily focused on subgroups of the labor market, such as inventors or entrepreneurs, studying the impact of exposure to occupations through parents (e.g., Bell et al., 2019), or the role of inherited of abilities (e.g., Lindquist et al., 2015; Aghion et al., 2017; Nicolaou et al., 2008), on the occupational choice of children.[6] We take a different approach. We study the whole range of occupations and ground our analysis in unique data and direct measures of how well individuals' skills match every occupation to quantify the macroeconomic implications of occupational persistence.

Second, our study relates to prior work that has documented how misallocation of talent and unequal opportunities during childhood can lead to lower educational attainment and earnings in adulthood (Chetty et al., 2016; Chetty and Hendren, 2018; Nakamura et al., 2021), 'lost Einsteins' (Bell et al., 2019), and reduced economic growth (Murphy et al., 1991; Hsieh et al., 2019). In light of this work, our results, finding talent not to be severely misallocated despite strong intergenerational persistence in occupations, may seem surprising at first sight. However, it is likely that through to the Scandinavian welfare state with tuition-free education and social security, Swedish men face lower barriers of entry into occupations than, e.g., women or immigrants due to factors such as differences in access to education (Goldin et al., 2006), labor market discrimination (Black and Strahan,

---

[6]More generally, we contribute to a vast literature studying the intergenerational correlation in a range of outcomes, including cognitive- and non-cognitive skills, educational attainment, labor market success, and health, see, e.g., Ermisch et al. (2012) for cross-country evidence and overview of literature. In addition to strong parent-child correlations in outcomes, prior work has documented strong sibling correlations, such as in education (e.g., Altmejd et al., 2021).

2001; Goldin, 2014), and social norms (Bertrand, 2011).[7] Still, there is a strong correlation in education, occupation, and income of Swedish fathers and sons. Our results imply that much of this intergenerational persistence stems from sorting on abilities, suggesting limited efficiency gains from policies aimed at reducing this persistence.

Third, our findings contribute to a literature documenting selection on abilities into education and employment in various settings. Exploiting a reform in compulsory schooling, Black et al. (2005b) conclude that high correlations in education of parents and children reflects selection rather than causation. Kirkeboen et al. (2016) find, using admission cutoffs and preferences of students, that choices of field of study are consistent with individuals choosing fields in which they have a comparative advantage. Fredriksson et al. (2018) study sorting of workers into jobs, documenting that selection is largely determined by skills. When combined with prior work documenting strong intergenerational correlation in both cognitive and non-cognitive skills (e.g. Black et al., 2009; Grönqvist et al., 2017; Björklund and Jäntti, 2012), our results are broadly consistent with the conclusion that intergenerational persistence in education, occupation, and income may reflect, to a large extent, transmission of abilities and sorting on skills.

This paper unfolds as follows. Section 3.2 describes the data we use. Section 3.3 documents patterns of occupational choice and intergenerational persistence. Section 3.4 extends the standard two-occupation Roy model with entry costs and discounts, and discusses the potential implications of these costs and discounts for occupational choice and intergenerational earnings mobility. In Section 3.5, we present the structure of the general equilibrium model of occupa-

[7] Unfortunately, our analysis is restricted to men, due to availability of data on skills from the military draft.

tional choice, and also describe how we predict potential productivities. We presents the results from the estimated model in Section 3.6. Using the structural model, we perform the counterfactual experiment in 3.7. Section 3.8 analysis quasi-experimental evidence, and then Section 3.9 concludes the paper. Additional background material is relegated to an appendix.

3.2 Data on Labor Market Outcomes

We use several data sets covering the Swedish population, both registery data and full-population census, extending back to 1960. Data on earnings and other labor market outcomes are obtained from tax records. Demographic and other background information, including data linking parents and their children, are obtained from administrative records. All of this data is compiled by Statistics Sweden and was made available to us through the servers of the Institute for Evaluation of Labor Market and Education Policies (IFAU).

At the core of our analysis lies the study of intergenerational relationships between the occupations of parents and children. For the period of 1960 to 1990, we use data from the Swedish Census (*Folk- och bostadsräkningen*), conducted by Statistics Sweden at five year intervals during this period. The census records both occupation and industry of the working age population. Starting in 1996, we use data from the wage statistics register (*Lönestrukturstatistiken*), which gathers data from employers about their employees every year. From this source, for each year, we have information on the occupations of workers in the public sector. However, only half of the private sector is sampled every year. Occupations are classified according to a Swedish version (SSYK-96) of the International Standard Classification of Occupations (ISCO) codes. Using cross-walks between versions of the classifications that we obtain from Statistics Sweden, we have a consistent classification of 113 3-digit ISCO-88 level occupations for the period 1960-2013. In 2013 the occupation classification scheme changed substantially. In order to maintain a consistent classification for parents and children, we end our sample period there. Appendix 3.A.1 provides details on the occupation classification and our cross-walks.

As our analysis is focused on the persistence of occupations and income across generations, we measure these when individuals are in their prime age. More precisely, for children in our sample, the prime-age occupation is defined as the modal occupation between ages of 30 and 40, i.e., the occupation observed in most years. If two occupations tie according to this criterion, we define the prime age occupation to be the one observed last in the age span. Income at prime age is then defined as total yearly labor earnings while working in the prime age occupation.[8] For parents, prime age occupation and income are defined in the same way, except we choose a higher age range of 45—55 with the aim of increasing the number of parent-child occupation observations. We restrict our sample to occupations with at least 1,000 men in order to avoid small cells, especially when measuring worker' skill-matches and predicted earnings in occupations, as we describe below. As a result, our final data set includes 696,016 father-son pairs in 91 different occupations.

[8] Appendix 3.A.6 shows average earnings by occupation.

3.3 Intergenerational Continuity of Occupation

In this section we document the tendency for intergenerational continuity of occupation—or occupational inheritance—in Sweden. First, in order to gauge how common it is for children to enter into the same occupations as their parents, we compute the share of sons and daughters, who follow their fathers and mothers into the same 3-digit occupation. Figure 3.1 reports the results for each of the four cases.[9] Occupational following is prevalent: Around 15 percent of women and twelve percent of men inherit one of their parents' occupations. Moreover, the figure reveals that following is divided along gender lines: sons are about three times as likely to follow their father as they are to follow their mother; similarly, daughters are more than four times as likely to follow their mothers as they are to follow their fathers.

3.3.1 Occupational Mobility

The simple statistic on how many children follow into the occupations of their parents has, however, a clear limitation: even if sons and daughters made their occupational choices independently, a certain fraction would, by coincidence, end up in the same occupation as their parents. Hence, the measure in Figure 3.1 is mechanically influenced by occupational sizes; in an economy with a few large occupations, following would be more common than in an economy with many smaller occupations.

[9] Figure 3.1 reports statistics based on the full sample. However, much of our analysis is based on the cohorts of men for whom we have data on skills. In that sample the share of sons that follow into their fathers occupation is similar to the full sample, or 8.6 percent compared to 8.8 percent.

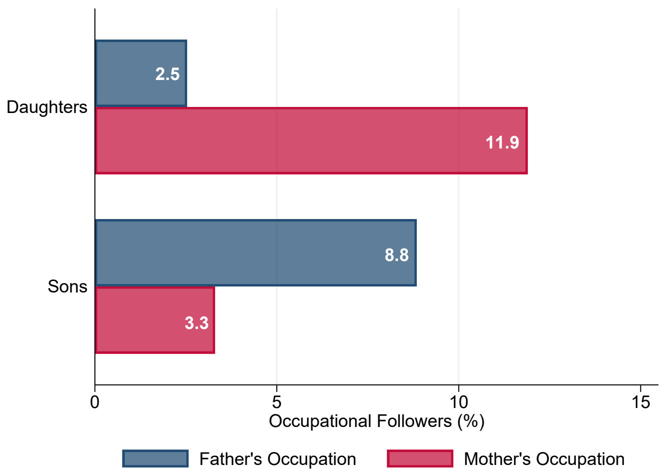

Figure 3.1: *Share of Followers*

Note: This figure shows, for sons and daughters, respectively, the fractions who choose same occupations as their mothers (red) and fathers (blue). The sample period is 1960-2013. Occupations are prime age occupations (see text).

To address this mechanical problem, we follow Rogoff (1953) and compute what we refer to as the *occupational mobility bias*, defined as:[10]

$$\mathsf{OMB}_{\mathsf{f},\mathsf{k}} = \frac{\mathsf{share}_{\mathsf{f},\mathsf{k},\mathsf{child}}}{\mathsf{share}_{\mathsf{k},\mathsf{child}}}$$

where f and k index the parent's and child's occupations, respectively. The occupational mobility bias for occupations f and k, $\mathsf{OMB}_{\mathsf{f},\mathsf{k}}$, is defined as the share of children with a father in occupation f who are observed in occupation k, $\mathsf{share}_{\mathsf{f},\mathsf{k},\mathsf{child}}$, relative to the fraction of children in occupation k, $\mathsf{share}_{\mathsf{k},\mathsf{child}}$. Intuitively, if occupations were assigned to children at random, then occupational mobility bias would be equal to one: the share of children in occupation k, conditional on having a father in f, would be the same as the share of all children in occupation k. If more children are found in occupation k with their parents in occupation f than would be expected under random assignment, the occupational mobility bias rises beyond one.[11]

Figure 3.2 summarizes the occupational mobility bias across all possible combinations, f and k, for fathers and sons.[12] The y-axis represents the father's occupation, while the x-axis represents the

[10] As discussed in Blau and Duncan (1967), in the sociology literature this ratio has been referred to as the "index of association" and "social distance mobility ratio".

[11] Our measure, OMB, compares the probability of observing a child in occupation k conditional on the father being in occupation f to the unconditional probability of observing a child in occupation k. Dal Bó et al. (2009) compute the probability of observing a father in occupation f conditional on a child being in occupation k and compare it to the unconditional probability of observing a father in occupation f. They refer to this measure as *dynastic bias*. By Bayes' rule, the two are mathematically equivalent.

[12] In our main analysis, we focus on the occupational choices of sons, due to data limitations on skills of women. For the occupational mobility matrix for mothers and daughters, see Figure 3.30 in Appendix 3.D.

son's occupation. Each row or column in the matrix is a specific three-digit occupational code in the Swedish SSYK-96 system, the vertical and horizontal lines partition the space into one-digit occupational categories.[13]

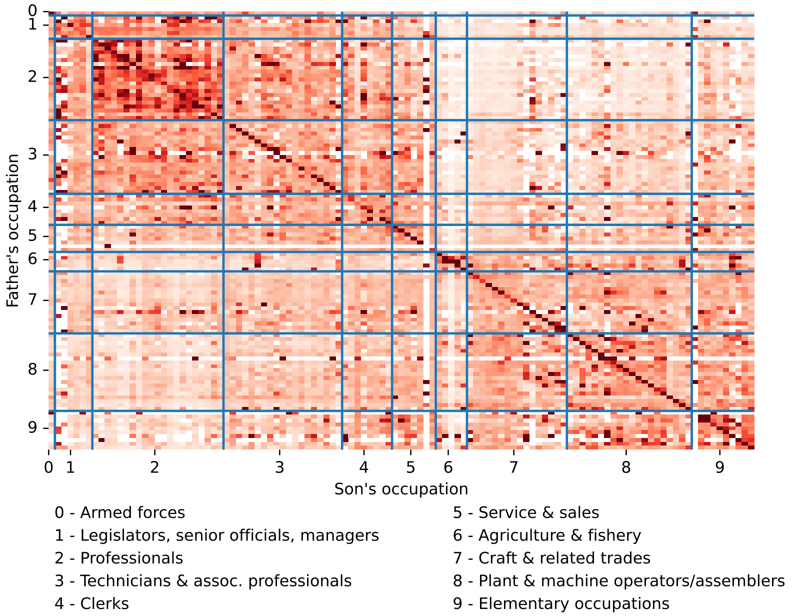

Figure 3.2: *Mobility Bias Across Occupations*

Note: This figure shows the mobility bias estimates across different occupations. The y-axis displays the father's occupation, the x-axis displays the son's occupation. On the x-axis, occupations are ordered according to their 3-digit code in the SSYK-96 classification system. The vertical and horizontal lines partition the space into 1-digit occupational categories. For the computation of the mobility bias, see the text. The sample period is 1960-2013.

The diagonal is clearly visible in the figure, implying that occu-

[13]For a list of occupational codes and descriptions, see Table 3.1 in Appendix 3.D.

pational following, i.e., sons choosing the same occupation as their fathers, is very common; occupational mobility bias is far in excess of unity. The weighted (unweighted) average of the bias along the diagonal is 6.04 (9.39), meaning that individuals with fathers in a given occupation are on average six times more likely to enter that same occupation, compared to an economy in which sons choose occupations at random. These findings are in line with previous studies that have documented substantial occupational mobility bias, e.g., on the US labor market (Rogoff, 1953; Blau and Duncan, 1967; Dal Bó et al., 2009). Perhaps surprisingly, the mobility bias is substantial across the whole range of occupations, and not concentrated to high- or low-paying occupations.[14]

Beyond the diagonal, there are clusters of occupational persistence. Especially among professionals, which include high-paying white collar occupations such as legal professionals, i.e., lawyers, and health professionals (except nursing), i.e., medical doctors and pharmacists, there is high mobility bias outside of, but close to, the diagonal. This implies that, e.g., while the children of health professionals are very likely to choose this occupation themselves, they are also more likely to stay within the broader occupational category than random assignment into occupations would predict.

Zooming in, Figure 3.3 shows the mobility bias for those sons who follow their fathers, i.e., the diagonal of the heatmap. The bias is highly heterogeneous across occupations, but almost always greater than one (note that the x-axis displays the bias in log-scale). We register the highest mobility bias among sons who choose agricultural professions, with values exceeding 100. The only profession for which

[14]Figure 3.3 plots a bar graph of the diagonal of the matrix in Figure 3.2.

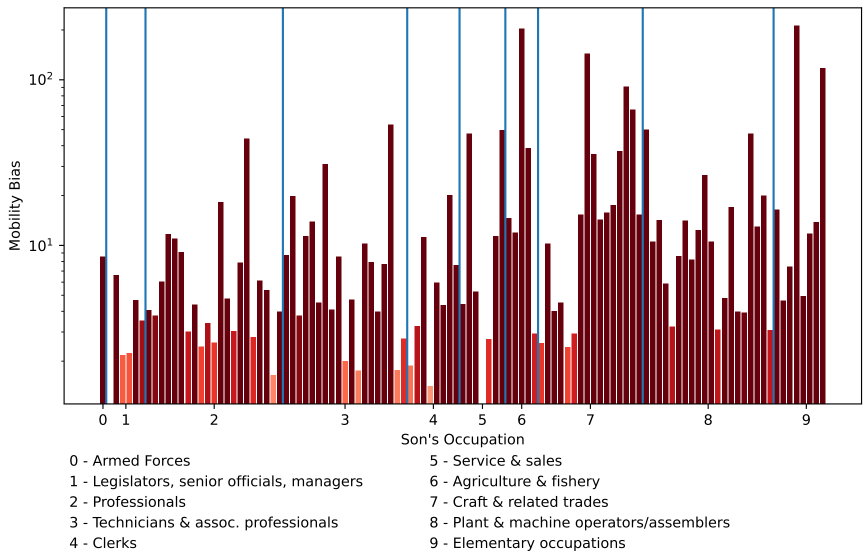

Figure 3.3: *Mobility Bias – Occupational Following*

Note: This figure shows a bar graph of mobility bias for children following their parents into the same occupation, i.e., f = k. The values are equivalent to those on the diagonal of Figure 3.2. On the x-axis, occupations are ordered according to their 3-digit code in the SSYK-96 classification system, the horizontal lines mark the borders of 1-digit occupational groups. The sample period is 1985-2013.

the mobility bias is smaller than one can be found among clerks.

Zooming out, the heatmap appears to be split into four sectors. The north-east and south-west corners show noticeably lower occupational mobility bias, while the south-east and north-west quadrants show noticeably more. Occupations up to and including service and sales can mostly be characterised as white-collar, i.e., police officers, lawyers, doctors, teachers, etc; while the occupations with one digit codes from six to nine are blue-collar occupations, i.e., fishermen, painters, and machine-operators. Sons are more likely to stay within these broad occupation types than random assignment would imply, and there is little movement across the two, as signified by bias below unity.

This relative immobility in intergenerational occupational choice has important implications for intergenerational earnings mobility more generally. Figure 3.4 plots the relationship between the fathers' and the sons' prime income ranks, similar to Chetty et al. (2014). The rankings are constructed within cohort-year cells. In our sample, the sons of fathers with very low income ranks also rank relatively low in their own income distribution, on average. The opposite is true for sons of fathers in the top of the income distribution. Hence, the graph exhibits a strong positive slope, implying that intergenerational mobility in prime age earnings is somewhat limited in Sweden. The correlation between father's and son's earnings ranks is 0.243.

To show the importance of occupational choices on earnings mobility, we assign every son in our sample the average earnings of his occupation. Thus, we control for intergenerational earnings mobility due to sons performing particularly well *within* an occupation and isolate the mobility driven only by their occupational choices. The red diamonds in Figure 3.4 display the result of this exercise. The relationships between the fathers income ranks and the sons income ranks are almost unchanged when restricting variation to *across* occupation-differences. Hence we conclude that the relationship between the fathers and sons income ranks is primarily driven by occupational choices, and that earnings differences within occupations have only small effects. Consequently, understanding the intergenerational persistence of occupational choices can help us shed light on intergenerational mobility more generally.

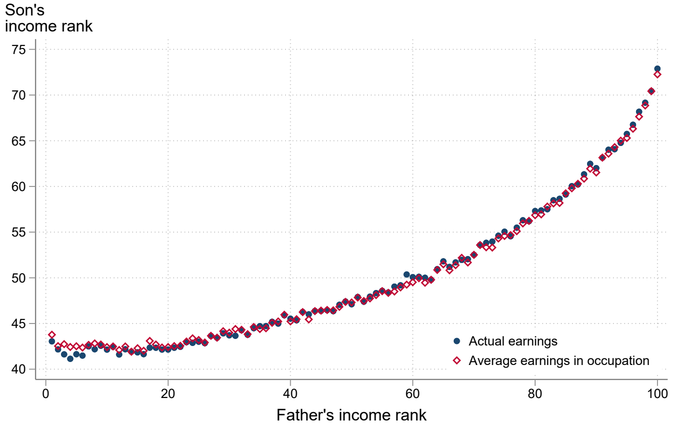

Figure 3.4: *Rank Rank Relation – Actual and Average Earnings*

Note: The figure shows the relationship between son's and their fathers' income ranks. Fathers are placed into 100 percentile bins. For each such bin, we calculate the average income rank of sons, which is then plotted on the y-axis. Fathers and sons are ranked within cohort-year cells. *Navy-colored dots* are based on observed earnings for the sons. *Red-colored diamonds* instead plot average income ranks, conditional on the income rank of fathers, when we instead of using individuals' actual earnings assign them the average income in their occupation. The sample period is 1985-2013.

3.4 A Basic Model of Occupational Choice

To study how skills and family background influence occupational choices and labor market outcomes, we build a Roy (1951) model that incorporates these factors. We build on and extend Roy models presented in Ohnsorge and Trefler (2007), Adão (2015), Nakamura et al. (2021), and, in particular, Mayer (2008). In the standard model, individuals are endowed with heterogeneous skills and choose between occupations where skills differ in their productivity. Importantly, we add two features to this setup. First, sons' skills partly depend on their fathers' skills, leading to intergenerational correlation in productivity across generations. Second, entering an occupation is costly and this cost may depend on the father's occupation. In this section, we present a simple partial equilibrium model in order to illustrate the mechanisms at play. In the subsequent section, we relax some simplifying assumptions in the model and extend it to a general equilibrium model that fits the Swedish economy. We estimate that model and use it to perform counterfactual experiments.

There are two occupations in the economy—hunting and fishing—in which an individual from family i and generation g can choose to work.[15] We use the generic index n to denote the occupations and denote fishing by F and hunting by H. Individuals live for two periods. In the first period, individuals from generation g are born as children of parents from generation $g-1$, learn their skills and choose an occupation. In their second period they are parents and inelastically supply one unit of labor to market work in their chosen occupation. This implies that in a given period only one generation is active in the labor market.

[15] We use g to denote both time and a generation, which consists of all individuals born in the same period, i.e., a birth cohort.

Occupations require an occupation-specific skill for workers to be productive.[16] Individuals are endowed with a bivariate skill vector $(Z_H^g(i), Z_F^g(i))$, where $Z_n^g(i)$ is the productivity of the individual from family i of generation g in occupation n. Each generation consists of a unit mass of individuals distributed across $\mathcal{Z}_F \times \mathcal{Z}_H$. We posit the distribution of Z_F^g in the population to be $F(Z_F)$ and the conditional distribution of Z_F^g to be $\{Z_F^g(i)|Z_H^g(i) = z\} \sim H(Z_F^g(i)|z)$.

We denote logarithms of variables in upper-case letters with a lower-case letter, i.e., $z_n^g(i) \equiv \log(Z_n^g(i))$. Children inherit skills from their parents with an error according to the following process of intergenerational persistence:

$$z_n^g(i) = \tau z_n^{g-1}(i) + (1-\tau)\varepsilon_n^g(i), \tag{3.1}$$

where τ governs the inheritability of skills. As $\tau \to 0$, children's abilities become independent of their parents' abilities, whereas $\tau \to 1$ implies that skills do not change from a parent to a child. The joint distribution of the skill innovations ε_n^g is assumed to be bivariate normal with mean $\mu_n = 0$ and variance $\sigma_n^2 = 1$. The correlation between the two skills is ρ. This leads to an ergodic distribution with mean $\bar{\mu}_n = 0$ and variance $\bar{\sigma}_n(\tau)$.

We assume, for simplicity, that labor is the only factor of production and firms produce using linear production functions:

$$Y_F = A_F L_F \quad \text{and} \quad Y_H = A_H L_H, \tag{3.2}$$

where

$$L_F = \int_{i \in \Gamma^F} Z_F^g(i)^{\beta_F} di, \quad L_H = \int_{i \in \Gamma^H} Z_H^g(i)^{\beta_H} di, \tag{3.3}$$

[16]We use the terms skills and abilities interchangeably to describe a fixed characteristic of a worker which governs their productivity.

Γ^n denotes the set of workers employed in occupation n, A_n represents aggregate productivity in sector n, and β_n represents the marginal return to productivity in sector n.[17] The labor markets for both occupations are perfectly competitive and firms operating in those markets take the prices of fish, P_F, and rabbits, P_H, as given. Here, we assume that prices are fixed, an assumption we relax when estimating the extended general-equilibrium model in the subsequent section. These assumptions imply that the wages per efficiency unit of labor in fishing and hunting, respectively, are given by

$$W_F = P_F A_F \quad \text{and} \quad W_H = P_H A_H \tag{3.4}$$

Earnings of worker i in occupation n is $Y_n(i) = W_n Z_n(i)^{\beta_n}$ and thus depends on the occupation's wage rate W_n, the number of efficiency units of labor the worker can supply $Z_n(i)$, and the marginal return to skills in the occupation, β_n. The logarithm of labor income is therefore given by

$$y_F^g(i) = w_F + \beta_F z_F^g(i) \quad \text{or} \tag{3.5a}$$

$$y_H^g(i) = w_H + \beta_H z_H^g(i), \tag{3.5b}$$

depending on whether the worker is a fisherman or a hunter, respectively. Without loss of generality, we assume that $\beta_F > \beta_H$. This echos an assumption made in the original Roy (1951) model, namely that "*rabbits are plentiful and stupid*" but the "*trout, on the other hand, are particularly wily and fight hard*".

[17]Our choice of modeling the marginal product of efficiency units using β_n follows Ohnsorge and Trefler (2007). Another common, and isomorphic, formulation is to assume that the variances of the intergenerational productivity innovations, ε_n^g differ across occupations (e.g., Sattinger, 1993).

Lastly, as children, individuals choose an occupation $k \in \{F, H\}$ that maximizes their utility in adulthood. Utility is log-linear and depends on three factors: earnings, y_n, an entry cost, m_n, and an entry-cost discounts, d_n. Entry costs are occupation specific, implying that any entrant has to incur them. Workers who follow their parents into the same occupation, however, secure a discount on the entry costs. Intuitively, this discount captures multiple forces which may make employment in their father's occupation more pleasant, convenient, or profitable: parents may facilitate better information about and access to necessary education (Lentz and Laband, 1989), may pass on knowledge (Laband and Lentz, 1983; Lentz and Laband, 1990; Laband and Lentz, 1992), provide contacts (Kramarz and Skans, 2014; Dal Bó et al., 2009), or simply bequest the family business to their children (Mocetti, 2016). Hence, utility is

$$u(i, g, n) = y_n^g(i) - m_n + d_n \mathbb{I}_{i^{g-1}, n = i^{g-1}, k}, \tag{3.6}$$

where $\mathbb{I}_{i^{g-1}, n = i^{g-1}, k}$ is an indicator function for having a parent in occupation n. The entry-cost discount acts as a pull factor for children with a parent in occupation n. If the discount is large, more children with parents in occupation n will follow, all else equal. For simplicity, we assume that parental discounts are zero for all generations $g < \underline{g}$. In what follows, we analyse how entry discounts in the model affect intergenerational mobility between generations $g - 1 = \underline{g} - 1$ and $g = \underline{g}$.

Figure 3.5 outlines the main mechanism in the model. It plots individuals' utilities in fishing (dark blue) and hunting (light blue) depending on their relative productivity in fishing compared to hunting, $s \equiv \beta_F z_F - \beta_H z_H$. It is useful to think of this as determining

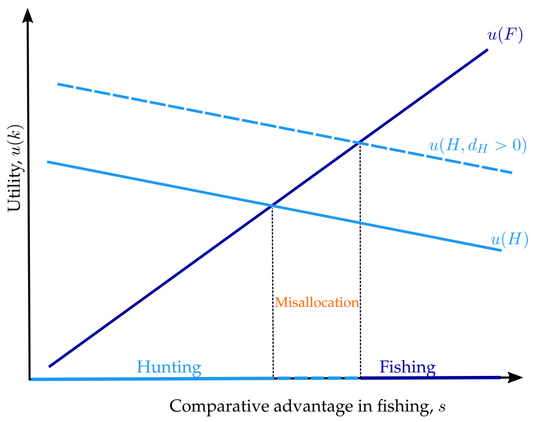

Figure 3.5: *Occupational Sorting by Comparative Advantage*

Note: The figure illustrates sorting into occupations based on comparative advantage and the effect of parental background on occupational choice. For simplicity, the figure illustrates the case where only sons of hunters receive a discount on the entry cost into hunting. This leads to increased entry of hunting sons into the hunting, despite them having comparative advantage in fishing, i.e. misallocation of talent. The case of discount on the entry cost into fishing is analogous.

an individual's *comparative advantage* in fishing, with the shorthand s referring to sorting. Similarly, $a \equiv \beta_H z_H$ measures a worker's *absolute advantage*. By rewriting equations (3.5a) and (3.5b) in terms of s and a, one can see that a change in a shifts y_F and y_H—and therefore $u(F)$ and $u(H)$—by the same amount.

Individuals with a large s are relatively more skilled as fishermen than hunters, i.e., have a comparative advantage in fishing, and choose to become fishermen. Given s, individuals who have a high a are highly productive in both occupations, i.e., have an absolute advantage in both fishing and hunting.[18] Furthermore, under the assumption that $\rho > 0$, those that become fishermen also tend to be skilled hunters, i.e., have a high absolute advantage in both occupations. Those that choose to become hunters, however, tend to have a low absolute advantage in both occupations, but a comparative advantage in hunting. Under $\rho < 0$ the reverse is true. In this section we assume $\rho > 0$, in line with the cross-sectional correlation in skills in the Swedish data. This assumption simplifies the discussion that follows on the model implications for intergenerational mobility. When extending this model and bringing it to data, we do not, however, need to make assumptions about skills or their correlation, as these are measured in the data.

Occupational choice in this model is directly influenced by occupational choices of parents. Figure 3.5 displays this influence on the occupational choices of children of hunters.[19] Having hunter parents shifts the line reflecting utility in hunting upwards, inducing more

[18]This can be seen from the definition of s: for a given s, a high $Z_H^{\beta_H}$ implies high $Z_F^{\beta_F}$.

[19]The case where children of fishermen receive a discount into fishing, not depicted, is analogous and would be represented with an upward shift of the dark blue line and an increase in the share of fishermen.

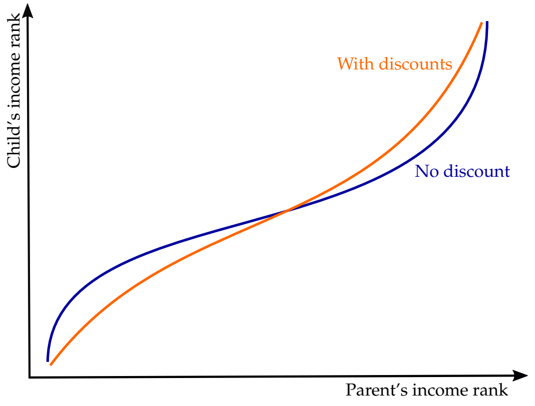

Figure 3.6: *Intergenerational Income Mobility*

Note: The figure presents the relationship between the income rank of children and their parents in the case of selection only on comparative advantage (dark blue) and with discounts on entry costs into parents' occupations (orange).

children to follow their parents into hunting. Absent parental discounts, however, these workers would have selected into fishing based on their comparative advantage. Therefore, parental discounts misallocate talent and distort efficiency.

Importantly, this model also allows us to study how parental influence on occupational choices can affect intergenerational mobility. In the model, as in the data, we measure intergenerational mobility by the relationship between the earnings rank of sons relative to other sons in generation g and the earnings rank of fathers within generation $g-1$. Figure 3.6 plots the rank-rank relationship in the model. In drawing this graph, we maintain our assumption that $\beta_F > \beta_H$ and in addition assume that $w_F > w_H$. That is, fishing can be thought of as the higher-paying occupation. In line with that assumption, we also assume that entry costs are larger in the fishing occupation,

$m_F > m_H$. The figure presents the rank-rank relationship for two cases: with and without discounts on entry costs based on parental background. The discounts lead some children of fishermen to choose fishing and some children of hunters to choose hunting, despite their comparative advantage lying in the other occupation. For children of fishermen, the discounts allow them to enter the higher-paying occupation, leading them to earn higher incomes than otherwise. For children of hunters, the discounts keep them in the lower-paying occupation, leading them to earn lower incomes than otherwise. Together the discounts decrease intergenerational income mobility, depicted as steepening the slope of the rank-rank relationship.

To summarize, the model provides two testable predictions. First, parental influence on children's occupational choices will increase the intergenerational persistence in occupations which reduces intergenerational income mobility. Second, parental influence distorts the efficient allocation of talent in the economy. The size of these effects will depend on the importance of parental influence relative to selection on skills in explaining the observed persistence in occupations of parents and children.

3.5 General Equilibrium Model of Occupational Choice

We now extend the basic two-occupation model from the previous section to a structural multi-occupation model that we can estimate using Swedish administrative data on skills and labor market outcomes. This model is able to replicate the important patterns of occupational choice observed in the data. We then use this model to perform counterfactual experiments to investigate how the different drivers of occupational following influence productivity and intergenerational income mobility. A central ingredient for the model estimation is a measure of how well individuals fit to different occupations. We measure this predicting the potential earnings of every individual in every occupation he could choose.

3.5.1 Potential Earnings

Data on Skills

We use a detailed measure of individuals' skills, utilizing scores from tests administered at military enlistment. These measures are available from the Swedish Military Archives from 1969. During our sample period, almost all men went through a draft procedure at age 18 or 19. The draft procedure consists of standardized tests that measure cognitive skills along four dimensions and a structured evaluation based on behavioral questions by a trained psychologist that evaluates individuals' personality traits (non-cognitive skills) along four dimensions. The cognitive skills are (1) *Logic-inductive ability* (fluid intelligence), (2) *Verbal comprehension* (crystallized intelligence), (3) *Spatial ability*, and (4) *Technical understanding*. The non-cognitive skills or personality traits are: (5) *Social maturity* (extroversion, having friends, taking responsibility), (6) *Intensity* (the capacity to activate oneself without external pressure, the intensity and frequency of free-time activities) (7) *Psychological energy* (perseverance, ability to fulfil plans, to remain focused), (8) *Emotional stability* (ability to control and channel nervousness, tolerance of stress, and disposition to anxiety). For further information about these measures, see Carlsted and Mårdberg (1993) and Mood et al. (2012). Previous work has documented the cognitive and non-cognitive test scores are correlated, but contain independent information about individuals' abilities and traits (Fredriksson et al., 2018).

Skill-Based Predictions of Potential Earnings and Occupational Fit

Our conceptual approach to measuring how well individuals fit to occupations based on their skills is motivated by the "task framework" (Autor et al., 2003; Gibbons and Waldman, 2004). According to this conceptual framework, occupations differ in tasks as well as skills required to perform these task. As individuals are heterogeneous in their skills, they differ in how productive they are in different occupations. Based on this, our presumption is that occupations differ in returns to skills. This is in line with results from prior work documenting heterogeneous returns to skills, e.g., higher returns to cognitive skills in occupations where such skills are a complement to technology (Katz and Murphy, 1992; Acemoglu and Autor, 2011; Hermo et al., 2022) and high returns of non-cognitive skills in occupations requiring significant interpersonal interactions (Deming, 2017; Edin et al., 2022). By extension, this implies that the skills of incumbent workers can be used to measure the skill returns and requirements in each occupation.[20]

Our empirical approach to measuring skill-based potential earnings is to first train a machine-learning algorithm using the combina-

[20]The nature of this approach is similar to that in Fredriksson et al. (2018), who use individuals' skills to measure mismatch in jobs. They assume, and empirically document, that workers sort into jobs based on skills of the incumbent workers in the job. They measure mismatch as the difference in an individual's skills and the average skills of incumbents in the same job. In contrast, our approach, as outlined below, is instead to estimate potential earnings in each occupation based on a flexible combination of all skills of incumbent workers in occupations, quantifying returns to skills in each occupation.

tion of skills and earnings of incumbants in each occupation and then predict potential earnings for all individual-occupation pairs. This procedure approximates an individual's productivity in each occupation. We also use a similar algorithm which predicts entry probabilities across occupations for each individual, which we use as a measure of occupational fit, i.e., match quality. Under the assumption that earnings reflect productivity, we base our predictions of entry probability—or occupational fit—on the skills of the most productive workers in each occupation, measured as workers in the highest quintile of the within-occupation earnings distribution. For earnings we instead use the whole distribution of earnings within an occupation to measure the productivity of different skills and skill compositions, exploiting that earnings are increasing in skills but differentially across occupations. In both cases, the training sample for the prediction is based on a sample that excludes individuals that follow their fathers into the same occupation, in order to avoid the influences of other characteristics than skills that may influence earnings and entry probability of followers.[21]

For our training and prediction, we use a random forest algorithm (Breiman, 2001), which constructs a multitude of decision trees along splits of skills and predicts an outcome by aggregating over the predictions of the individual trees. The algorithm then minimizes the root mean squared error (RMSE) between predictions and observed realizations for multiple training samples. The usefulness of this method is its flexibility, as skills are likely to be required in various degrees and interactions across different occupations. In this sense, the random forest is superior to, e.g., a simple regression of individual earn-

[21] In practice, this restriction has limited quantitative influence on the predictions, as those based on sample that excludes vs. includes followers have a correlation of about 0.98.

Figure 3.7: *Actual and Predicted Earnings*

(a) *Within-Occupation Rank of Earnings*

(b) *Rank of Earnings*

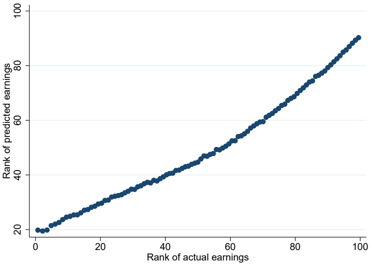

Note: This figure plots the relationship between predicted and actual earnings, presented in ranks for comparability across occupations. Panel (a) plots the average within-occupation rank of predicted earnings for individuals in a specific bin of actual within-occupation earnings. Panel (b) plots the average rank of predicted earnings across ranks of actual earnings. Earnings are predicted by a Random Forest algorithm using individual skills as inputs. Occupational followers are excluded from the estimation.

ings on skills, which would impose linearity on the relationship and not allow for exhaustive possibilities of interactions of skills. In practice, for each occupation, we predict individual residualized earnings in logarithms, that is, residuals from a regression on age, year and occupation fixed effects. For our model estimation and analysis, we convert the predicted residuals into values in Swedish Kronor (SEK), using the estimated fixed effects, normalizing earnings by time and age.[22] Appendix 3.B provides a detailed description of the estimation procedure.

We find that cognitive and non-cognitive skills have substantial

[22]For comparability of earnings across individuals within occupation, we normalize earnings to that at age 40 in a period, where we split our sample into six periods (two every decade) when individuals are in their prime age.

predictive power of entry probability and earnings within occupations.[23] Figure 3.7 shows the relationship between the earnings predictions obtained from our random-forest algorithm and actual earnings of incumbents. In Figure 3.7a we plot the within-occupation rank of predicted earnings against the rank of actual earnings, across all occupations. The figure displays a strong positive correlation between the skill-based predictions of earnings and actual earnings.[24] In Appendix Figure 3.27 we plot the histogram of R^2 from the random-forest predictions, by occupation, which average to 9.3. In Figure 3.7b, we plot the relationship between predicted and actual earnings, presented as ranks within birth cohort and year. Similar to panel (a), the graph depicts a strong positive correlation between our earnings predictions and actual earnings. Overall, we find strong signs of our approach being able to credibly map skills to earnings.

As described above, the hypothesis underlying our approach is that skills are differently productive in different occupations. To empirically evaluate this hypothesis, we document the relative impor-

[23] Appendix Figure 3.32 plots the histogram of predicted probabilities of occupational entry. The figure documents a dominantly higher probability for incumbents, including those not at the top of the earnings distribution and therefore not used in training the algorithm. This implies that lower-income workers select into their occupations on a similar set of skills as do workers at the top of the earnings distribution. This supports our approach for predicting earnings, which uses incumbents across the whole earnings distribution.

[24] As the figure documents, while we are able to obtain a qualitatively good prediction of earnings, it is quantitatively imperfect, as displayed by the considerably smaller range of the predicted earnings than the range of their empirical counterpart. This is expected, as the prediction is solely based on skills, while actual earnings reflect a range of other factors. In the structural model, this introduces a bias which pushes the economy towards greater intergenerational mobility compared to the data. In an extreme case in which predicted earnings were random, we would expect the intergenerational rank-rank relationship produced by the model to be flat. However, as the results in Section 3.6 indicate, this is not the case; the model's baseline results on intergenerational mobility are very close to the data.

Figure 3.8: *Factor Importance*

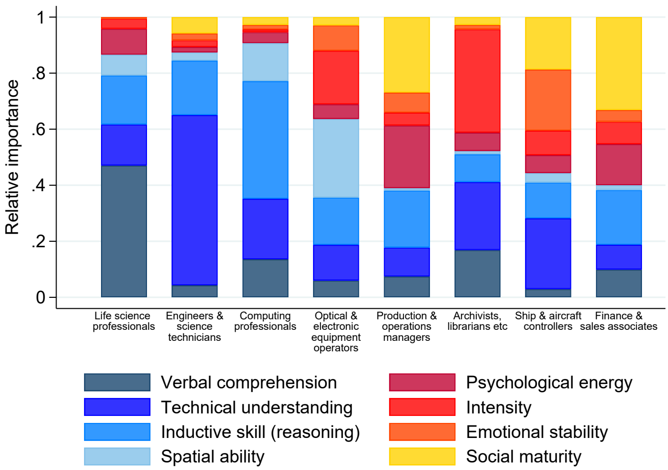

Note: This figure shows the relative importance of our eight skill measures in predicting incomes across occupations. The selected occupations are those in which each of the eight skills contributes the most to the overall prediction of income (see text for details). Occupations are ordered along the x-axis by cognitive (left) and non-cognitive (right) skills. Relative importance measures the contribution of a split along a given skill to the prediction.

tance of each of the eight skills in predicting earnings in occupations. Figure 3.8 plots a measure of relative importance that is based on the contribution splits along the dimension of each skill to the overall prediction of income. The figure plots this measure for eight different occupations, selected and ordered based on the relative importance of each skill. It shows that occupations differ substantially in the relative importance of skills, but also that a variety of skills are productive in each occupation. For life science professionals, and engineers and computing professionals, cognitive skills (four shades of blue) are by far the most important in predicting earnings, contributing about 90 percent to the earnings prediction. In contrast, for production and operations managers, and finance and sales associates, non-cognitive skills are relatively more important and contribute about 60 percent.[25]

The evidence above supports the notion that skills are differentially productive across occupations, showing that occupations differ greatly in which skills are important in predicting income in a given occupation. In order to further evaluate this conclusion, we compare our measure of skill requirements based on incumbents in an occupation to a measure of skills required to solve the tasks performed in occupations. More precisely, following Macaluso (2017), we use the O*Net task-data to measure skill-distance between occupations and compare this to the skill distance implied by our predictions. As documented in Appendix 3.A.4, we find the two measures to be similar.

A valid concern is that the skills that we measure, and consequently predicted earnings and occupational fit, might partly be a result of upbringing. If this is the case, then we would underestimate how much background factors affect outcomes, such as occupational

[25]In addition to this evidence on the importance of skills across occupations, Appendix 3.A.5 documents that the average level of skills remains stable over time.

choice and earnings. We do two things to address this concern. First, in Appendix 3.A.2, we utilize that we have measures of skills in the early teens for a subsample of sons. We show that the relationship between skills and background, measured as either the father's skill rank or income rank, is strikingly similar in the early teens as compared to in the late teens, indicating that, if upbringing can affect test scores, it is not correlated with background. Second, in Appendix 3.A.3, we utilize brother-pairs. We find that the probability of entering the father's occupation increases with occupational fit, and that the magnitude of this relationship is only marginally affected when comparing brother-pairs, indicating that the relationship is *not* driven by upbringing.

3.5.2 Model Structure

Every individual is endowed with a Q-dimensional vector of skills $x = \{x_1, x_2, \ldots, x_Q\}$, where x_q measures the ability in dimension q. Individuals apply those skills to production in their chosen occupation according to an occupation-specific production function that takes their skills as inputs: $Z(x, n) = V_n(x)$. As in the basic model, individuals supply labor inelastically to the market where perfectly-competitive firms operate. Firms use labor as the only factor of production in a linear production function, as described by (3.2), and pay workers their marginal products. Earnings of worker i with skills x in occupation n is therefore $Y(x(i), n) = P_n A_n Z(x(i), n)$.

Individuals choose the occupation which maximizes their utility. We modify the utility function (3.6) in two ways. First, instead of assuming that utility is linear in earnings, we posit that individuals derive felicity according to the function $g(c_1, \ldots, c_N)$, where c_n represents consumption of goods produced by occupation n. They are subject to a budget constraint of the following form:

$$\mathbb{I}_1 Y(x(i),1) + ... + \mathbb{I}_N Y(x(i),N) = Y(x(i),k) = \sum_{n=1}^{N} P_n c_n(i) \quad (3.7)$$

where P_n is the price of goods produced in occupation n. The left-hand side of the equation represents worker's income, depending on his choice of occupation k, noted with the indicator $\mathbb{I}_k$. This formulation allows us, in general equilibrium, to derive demand functions for different goods, given a price vector.

The second modification assumes that utility is influenced by preferences over occupations. We model this with preference shocks $\varepsilon_n(i)$ which are i.i.d. across workers and occupations. These preference shocks serve two purposes: (i) they lead individuals with the same skill set x and father's occupation f to choose different occupations, which helps us match the empirical occupation distribution, similar to an approach common in spatial sorting (e.g., Diamond and Gaubert, 2021), and (ii) they convert the decision problem from one of discrete choice to one with nondegenerate choice probabilities (McFadden, 1974).[26]

As before, choosing an occupation n is associated with a utility cost, b_n^f, which consists of a general utility cost and a possible discount on entering the occupation n, which depends on father's occupation, as we describe in more detail below. In the next section, we

[26]To facilitate this, we assume that there is a measure $M_{x,n} \in \mathbb{R}_+$ of individuals in each cell of the skill-occupation distribution. In the data, naturally, we observe a discrete number $\delta_{x,n}$ of individuals in a skill-occupation cell, each of whom can only choose to work in a single occupation. With the assumption of a measure $M_{x,n} = \delta_{x,n}$ in each cell, we are able to smooth the problem, splitting each discrete worker into an infinity of workers. Shares of the measure can then be assigned to different occupations.

estimate these costs and discounts such that they match prominent features of the father-son occupational transition matrix in the data.

As before, the model is static with a single period. At the start of the period, each individual i with a father in occupation f takes prices $\{P_n\}_{n=1}^N$ and entry costs across occupations $\{b_n^f\}_{n=1}^{n=N}$ as given and solves the problem by backwards induction. First, he maximizes his consumption utility $g(\cdot)$ subject to the budget constraint, given his skill set x and every possible occupation n he can choose. This yields the indirect consumption utility function $h(n,s) = (c_1^\star(n,x),\ldots,c_N^\star(n,x))$. To choose their optimal occupation, individuals maximize their utility $u(f,k,x)$, subject to the cost vector they face and their individual preference shocks. We can now define the equilibrium in the economy just outlined.

An equilibrium in this economy is a set of prices $\{P_n\}_{n=1}^N$, such that, given costs $\{b_n^f\}_{n=1,f=1}^{n=N,f=N}$,

- Supply equals demand in all occupations n:

$$C_n = A_n Z_n \quad \forall n$$
$$\text{where } C_n = \int c_n(i)\, di, \text{ and } Z_n = \int_{i\in\Gamma^n} Z(x(i),n)\, di$$

  where Γ^n is the set of workers who choose to enter occupation n.

- Workers choose occupations optimally and maximize their utility.

3.5.3 Estimation

When estimating the model, we set the function $g(\cdot)$ to be a Cobb-Douglas aggregator across all the goods produced by different occupations:

$$g(c_1, ..., c_N) = \prod_n c_n^{\alpha_n} \quad \text{with} \sum_{n=1}^{N} \alpha_n = 1 \tag{3.8}$$

which gives the associated price index $P = \prod_n \left(\frac{P_n}{\alpha_n}\right)^{\alpha_n}$. This formulation is convenient, as it implies, combined with the budget constraint (3.7), that the optimal expenditure shares on each of the different products in the economy are governed by their respective α coefficients:

$$\alpha_n = \frac{E_n}{E}, \ \forall n \in N \tag{3.9}$$

where $E_n = P_n C_n$ and $E = \sum_{n=1}^{N} E_n$. Thus, the indirect consumption utility function, given an occupational choice k and prices, is a linear function of income $Y(x, k)$. We postulate that utility from consumption, costs associated with occupational choice, and taste shocks, are additively separable. Hence the total utility obtained by an individual with skills x and a father in occupation f who chooses occupation k is

$$u(f, k, x, i) = h(k, x) - b_k^f + \varepsilon_k(i) \tag{3.10}$$

The taste shocks $\varepsilon_n(i)$ are i.i.d. across workers and occupations. They are distributed according to a Type I Extreme Value distribution with parameter κ.[27]

As outlined in section 3.3.1, a striking feature in the data is the fact that a disproportionately large fraction of individuals choose ei-

[27]The PDF of the Type I EV distribution is $c(\varepsilon) = \kappa e^{-\kappa\varepsilon} e^{-e^{\kappa}\epsilon}$, and its CDF is $C(\varepsilon) = e^{-e^{-\kappa\varepsilon}}$. It can be shown that the mass of workers ψ_n who choose

ther the same occupation as their fathers, or an occupation that is similar. To account for this in the model, we let the costs $\{b^f_n\}_{n=1,f=1}^{n=N,f=N}$ vary with the occupation of the father in the following way. First, all individuals who enter occupation n pay an entry cost of m_n. These costs are the same for all sons, no matter which occupation their fathers hold. Additionally, we assume that, depending on his father's occupation, a son enjoys reductions in occupational entry costs. These reductions are additively separable and come in three stages: sons can (i) choose the same occupational type (blue collar/white collar), (ii) choose the same broad occupational category (one-digit occupational group), or (iii) choose to follow their father into the same occupation. A son who chooses to be a doctor and has a father working as a motor vehicle driver, therefore, enjoys no reductions, facing only the entry cost m_n. If his father was a doctor, however, he would receive all three reductions. Intuitively, the discounts capture multiple forces which may make employment in their father's occupation, or a similar occupation, more pleasant, convenient or profitable: parents may facilitate better access to education (Lentz and Laband, 1989), may pass on knowledge (Laband and Lentz, 1983; Lentz and Laband, 1990; Laband and Lentz, 1992)), provide contacts (Kramarz and Skans, 2014; Dal Bó et al., 2009), or simply bequest the family business to their children (Mocetti, 2016).

Let $G_n \in \{1, 2\}$ be the type of occupation n, i.e., white collar or blue collar. Furthermore, let $g_n \in \{0, ..., 9\}$, be the broad, one digit

---

occupation n is

$$\psi_n = \Pr(\text{argmax}_n \quad u(f, k, x) = n) \tag{3.11a}$$

$$= \frac{e^{\kappa u(f,k,x)}}{\sum_n e^{\kappa u(f,n,x)}} \tag{3.11b}$$

occupational category of occupation n. The cost that an individual with a father in occupation f has to pay to enter occupation n is then given by

$$b_n^f = m_n - \mathbb{I}_{G_f=G_n} d_{1,G_n} - \mathbb{I}_{g_f=g_n} d_{2,g_n} - \mathbb{I}_{f=n} d_{3,n} \qquad (3.12)$$

where d_{1,G_k} is the discount for individuals choosing the same type of occupation as their father, d_{2,g_k} is the discount for individuals choosing same broad occupational category as their father and $d_{3,k}$ is the discount for individuals choosing the same occupation as their father. Note that, in our case, there are two distinct d_{1,G_k}, one for white-collar and one for blue-collar, ten distinct d_{1,g_k}, and 91 distinct $d_{3,k}$.

Potential Earnings Across Occupations

A central input into our model estimation are individuals' skills and their resulting productivity across occupations, i.e., $Y(x(i), n)$. To make this connection, we first interpret the earnings that we observe in a person's prime age occupation as a measure of their individual productivity.[28] Moreover, without loss of generality, we normalize $P_n = 1 \ \forall n$, which implies that earnings within an occupation is equal to the number of units or services produced: a legal professional who earns 500,000 SEK per year is assumed to produce 500,000 units of legal services. The normalization has no effect on relative predicted

[28]For that to be true, we assume that all individuals work the same number of hours in a given occupation. If hours differ systematically across occupations, this will be absorbed in occupation fixed effects (see below). If individuals within the same occupation worked different numbers of hours, depending on their real productivity, our assumption would be problematic. Our model's implications for aggregate output would still hold true, but they could not easily be converted into implications for productivity.

earnings across individuals within occupations, which importantly is what matters for our results. Then, using the earnings predictions based on skills presented in Section 3.5.1, we obtain a productivity for every individual across all occupations.

Costs and Discounts

Given the aforementioned earnings predictions, in the model, we jointly estimate the costs $m = \{m_n\}_{n=1}^N$ and discounts $d_1 = \{d_{1,G_n}\}_{G_n=1}^2$, $d_2 = \{d_{2,g_n}\}_{g_n=1}^{10}$, and $d_3 = \{d_{3,n}\}_{n=1}^N$ to match a set of data moments within each of six data subperiods, in order to keep earnings values comparable over time. First, we target the shares of individuals in each of the N occupations. In the data, we measure this share as the number of sons observed in occupation n divided by the total number of all sons. These moments pin down the entry costs, m. To estimate the discounts d_1, we target (i) the share of individuals who have a father in a white collar occupation and choose a white collar occupation, and (ii) the share who have a father in a blue collar occupation and choose a blue collar occupation. Similarly, for the discounts in d_2, we target the shares of sons who choose an occupation that is within the same broad group of occupations as the father's occupation. Lastly, for the discounts for following into the same occupation as the father, d_3, we, for each occupation, target the share of sons who choose the same occupation as their father. We normalize the entry costs into the armed forces, the following discount for white-collar occupations and the follower discount children with a father in the armed forces to zero.[29] To calibrate the parameter κ, which governs the variance

[29] In Appendix 3.C we describe how we find initial guesses for the respective entry costs and discounts.

of preference shocks, we target the ratio of earnings at the 75th and 25th percentile.

The model is able to closely replicate the targeted moments: the share of sons who have fathers in white (blue) collar occupations and choose a white (blue) collar occupation themselves is 68.70 (59.86) percent in the data and 68.71 (59.85) percent in the model. Furthermore, the shares of sons who have an occupation in the same broad one-digit group as their father is reported in Figure 3.29 in Appendix 3.D. Again, the model's moments are very close to those measured in the data.

Figure 3.9 shows the comparison between the other model and data moments. The left panel displays the occupation shares in the model and the data, which pin down the occupation entry costs in the model. The largest difference between the two appears in the second digit 6 occupation, Animal producers and related workers, where the model over-predicts entry by 0.06 percentage points. On average, however, the difference between model results and targets, in absolute values, is close to zero. The right panel of Figure 3.9 shows the share of sons who follow their fathers, across all occupations. Here, too, the model comes very close to matching the targeted moments.

The model also does well along several other dimensions. It very closely reproduces the rank-rank correlation between fathers and sons previously shown in Figure 3.4. We show the model-generated relationship in Figure 3.31 in Appendix 3.D. Since the targeted statistics only account for a small share of all occupational choices, and intergenerational income mobility was not explicitly targeted, reproducing this graph is a success of the model. Additionally, the model matches the occupational choices of non-followers. Figure 3.10 shows the shares of children who choose four different occupations, sorted by

Figure 3.9: *Model Fit*

(a) *Occupation Density – Baseline Model and Data* ***(b)*** *Following – Baseline Model and Data*

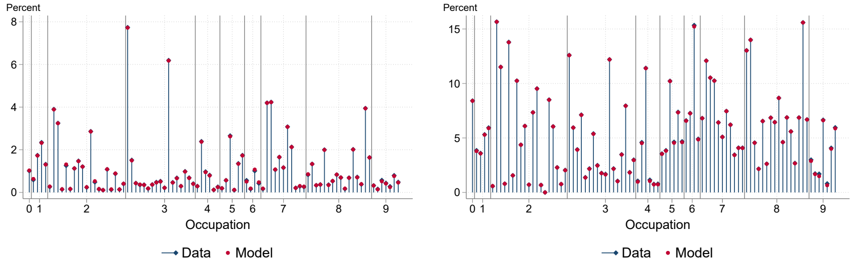

Note: The *Left Panel* shows the fraction of sons who choose each occupation. The blue diamonds represent this fraction for the pooled dataset, the red circles report results for the baseline model. The *Right Panel* shows, by occupation, the fraction of fathers whose child follows them into the same occupation. The blue diamonds represent this fraction for the pooled dataset, the red circles report results for the baseline model. On the x-axis, occupations are ordered according to their 3-digit code in the SSYK-96 classification system, the horizontal lines mark the borders of 1-digit occupational groups. The sample period is 1985-2013.

their fathers' income ranks. Importantly, followers, a fraction which was explicitly targeted in our calibration, are excluded from these graphs. The data shows that individuals born to fathers at the top of the income distribution are close to three times more likely to become health or legal professionals than sons born to fathers at the low end of the income distribution. Conversely, the children of low-earning fathers are much more likely to choose to become cleaners or mechanics than children of high-earning fathers.

(a) *Health Professionals (Except Nurses)*

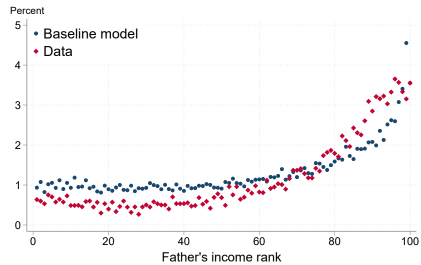

(b) *Legal Professionals*

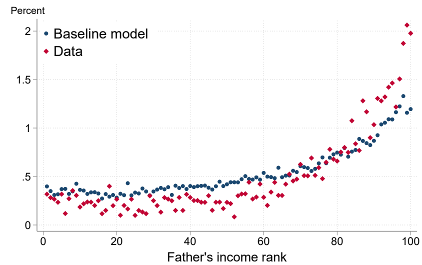

(c) *Helpers and Cleaners*

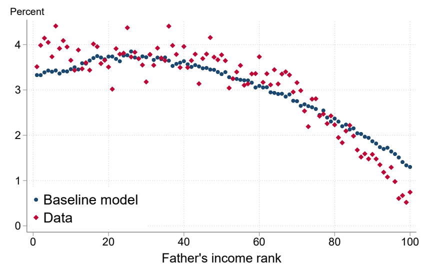

(d) *Machinery Mechanics and Fitters*

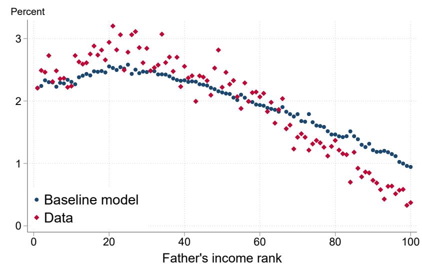

Figure 3.10: *Occupational Choice by Father's Income Rank*

Note: These figures plot the shares of individuals who choose four different occupations, depending on their fathers' income ranks. All figures exclude sons who choose the same occupation as their father, i.e., occupational followers. The blue dots represent the shares in the data; the red diamonds represent the shares in the calibrated baseline model. The *Top Left Panel* plots the share of sons who become health professionals, the *Top Right Panel* plots the share of sons who choose to become legal professionals, the *Bottom Left Panel* plots the share of sons who become helpers and cleaners, the *Bottom Right Panel* plots the share of sons who become mechanics and fitters. The sample period is 1985-2013.

3.6 Estimation Results

The left panel of Figure 3.11 displays the model implied costs of entering different occupations (blue), as well as the entry costs faced by individuals who can choose to follow their parent into the same occupation (red), i.e., each cost including all available discounts. We convert the utility costs and discounts into their respective monetary values.[30] Recall that we normalize the entry cost for the military professions (the leftmost occupation) to zero. The graph shows strong heterogeneity in entry costs. Towards the low end of the occupation spectrum, entry costs are high: becoming a director or chief executive, according to our model, carries the highest utility cost: the equivalent of almost 400,000 SEK *more* than entering a military profession. However, moving up through the occupational codes, entry costs fall below zero, relative to entry into military occupations. Our model estimates imply that the lowest entry cost is associated with animal producers and related workers.

The model implied discounts are most often relatively small.[31] The biggest difference between the utility entry costs faced by outsiders (blue diamonds) and followers (red circles) can be found in the first digit 9 occupation, helpers and cleaners. For sons who have a father in this occupation, following becomes very attractive. The same is true for pilots (the fourth digit 3 occupation), directors and chief

[30]Since, under the Cobb-Douglas assumption for $g(\cdot)$, income has a linear effect on income, we can map the utility cost of choosing an occupation into income by multiplying it with the price index $\prod_n \left(\frac{P_n}{\alpha_n}\right)^{\alpha_n}$.

[31]As can be seen in Figure 3.11, the estimated discounts are sometimes of "the wrong sign", indicating the followers pay an extra utility cost, and large in absolute values. The reason for this result is that the share of followers in these occupations are very low, and that the model, which uses taste shocks, requires an occupation to be *very* unattractive to generate very low choice probabilities for that occupation.

executives (the first digit 1 occupation), and doctors (the fifth digit 2 occupation). Anecdotally, this group of occupations makes clear that the discounts potentially capture very different types of exposures: Drycleaning businesses may be handed down from father to son, success as a chief executive likely depends on contacts and connections, and there may be significant informational frictions to becoming a pilot, which a father in the same occupation can reduce.

In some occupations entry discounts are negative, i.e., followers face a *larger* cost than non-followers. To replicate cases of little to no following in the data, the model requires such disincentives. This force is strongest among social science and linguistic professionals, the 15th digit 2 occupation.

To put the absolute values of the costs into perspective, the right panel of Figure 3.11 shows how *net* earnings are affected by the entry costs. The blue diamonds display the average potential earnings in each occupation n, net of the entry costs, for everyone in the economy, relative to earnings in the military:

$$\mathsf{earn}_n^{noc} = \frac{100}{N}\left(\sum_{i=1}^{N} \frac{\mathsf{earn}_{in} - \mathsf{cost}_n}{\mathsf{earn}_{i1}} - 1\right) \tag{3.13}$$

Hence, a value of -13 percent, as in the case for pilots (the fourth digit 3 occupation), implies that the average person in the economy would earn 13 percent less as a pilot than they would as a military professional, after taking entry costs into account. According to this metric, entering as a director or chief executive (first digit 1 occupation) appears to be the least attractive choice; individuals would have to give up more than one third of their prospective prime age earnings as military professionals to enter the occupation.

However, when we focus on individuals who actually choose each

Figure 3.11: *Model-Implied Costs*

(a) *Model Implied Entry Costs* ***(b)*** *Relative Earnings Net of Costs*

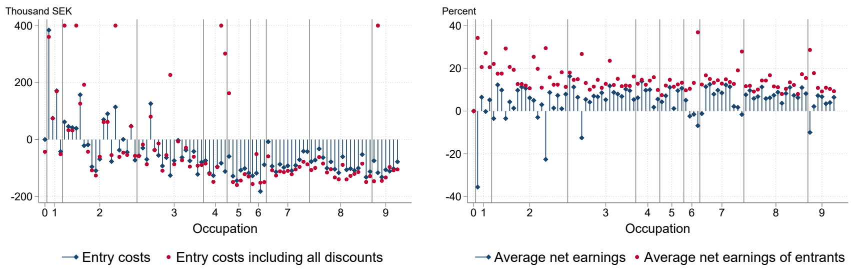

Note: The *Left Panel* shows the model implied entry costs in SEK (blue diamonds) and the costs for individuals following their father into the same occupation (red circles), i.e., the entry costs including all discounts. Estimated entry costs and discounts are period-and-occupation specific. In the current graph, we present averages, where each entry costs, and entry cost including all discounts, respectively, is weighted in proportion to the number of fathers in each occupation in each year. To make the graph more readable, entry costs including all discounts exceeding 400,000 SEK are top-coded. The *Right Panel* displays the earnings net of entry costs for everyone in the model, relative to military professionals (blue diamonds, see text for details) and the average earnings net of entry costs for individuals who choose each occupation, relative to military professionals (red circles). See text for more details. On the x-axis, occupations are ordered according to their 3-digit code in the SSYK-96 classification system, the horizontal lines mark the borders of 1-digit occupational groups.

occupation (red circles), earnings net of costs are much higher.[32] Those who choose to enter as a CEO actually earn 34 percent more,

[32]The estimated entry probabilities are strictly smaller than one, meaning that individuals are not allocated to just one occupation. We can, however, use these probabilities as weights, to calculate the average earnings net of entry costs for entrants: $earn_n^{enoc} = \sum_{i=1}^{N} \omega_{in} 100 * \left(\frac{earn_{in} - cost_n}{earn_{i1}} - 1\right)$ where $\omega_{in} = \frac{\eta_{in}}{\sum_{i=1}^{N} \eta_{in}}$. The weight ω_{in} represents individual i´s weight in occupation n, and η_{in} the estimated probability of entering. Naturally, $\sum_{n=1}^{N} \omega_{in} = 1$.

net of entry costs, than they would as a military professional. This implies that they possess a very particular set of skills fit for this occupation. The same is true, albeit to a lesser extent, in all occupations: entrants' predicted earnings, net of costs, are higher than for the average.

3.6.1 Interpreting the Cost Vector

In order to get a better understanding of what the estimated cost vector might be capturing, we relate it to time costs of entering an occupation. For this exercise, we utilize data from the Occupational Outlook Handbook of 2020.[33] In it, the BLS reports the typical education needed for entry into an occupation, as well as the typical work experience in related occupations (in years) that is required. Both of these measures are proxies for the time cost, and, hence, the utility cost, required to enter an occupation. For this reason, a positive correlation between these statistics and the model implied costs will serve as an indication that the model, together with our earnings predictions, can speak to occupational choices and their implications.

The educational requirement is split into eight categories: no formal educational credential, high school diploma or equivalent, some college (no degree), postsecondary non-degree award, associate's degree, Bachelor's degree, Master's degree, and doctoral or professional degree. We create a categorical variable that takes values 0 through 7 in the aforementioned order. Work experience is reported in three categories: none, less than five years and more than five years. Again, we assign categorical values from zero to two to each category.[34]

---

[33] Source: https://www.bls.gov/emp/tables/occupational-projections-and-characteristics.htm

[34] We map these statistics into the Swedish SSYK96 classification, as outlined

(a) *Model Costs and Educational Requirements*

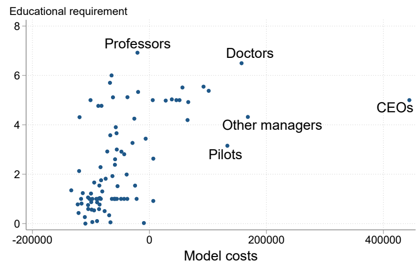

(b) *Model Costs and Usual Work Experience*

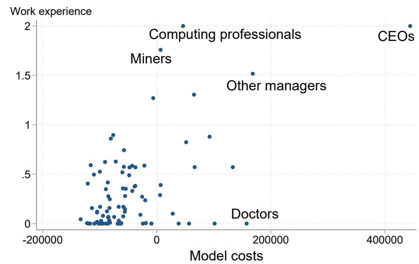

Figure 3.12: *Model Cost and Occupation Entry Requirements*

Note: The *Left Panel* plots the relationship between the entry costs estimated in the model (x-axis) and the educational requirements (y-axis), for different occupations. The educational requirement is coded as a categorical variable between 0 and 7 (see text). The *Right Panel* plots the relationship between the entry costs estimated in the model (x-axis) and the work experience in other occupations required for entry into an occupation (y-axis). The work experience is coded as a categorical variable between 0 and 2. Both educational and work experience requirements are obtained from the BLS' Occupational Outlook Handbook for 2020.

Figure 3.12 shows a scatterplot relating the model-estimated entry cost vector (relative to entry into the armed forces) to the educational requirements (left panel) and work experience (right panel) necessary for different occupations. In both cases, the costs estimated in our model calibration are strongly positively related to the data.

Our calibration implies that the highest utility costs accrue to CEOs, Pilots, managers and medical professionals. All of these professions either have high educational requirements (health professionals), or expect a lot of work-experience (managers and CEOs).[35]

in section 3.A.4.

[35]Additionally, we consider the role of occupational prestige or amenities in affecting the entry costs (Boar and Lashkari, 2021). In sociology, such prestige is often measured by the so-called Treiman scale (Treiman, 1977). Unfortunately, this

3.7 Counterfactual Analysis

3.7.1 Effects on Occupational Choice and Following

For our main counterfactual, we quantify the impact of entry-cost discounts, by setting them all to zero. This reduces the son's utility of choosing the same occupation as his father's, increasing the relative attractiveness of all other occupations. We then solve the model again, holding entry costs m_n fixed and letting the prices P_n adjust to clear the market.

Figure 3.13 shows how this change affects the share of sons who follow their father, out of all sons born into each occupation. The blue diamonds plot this fraction for the baseline specification, reproducing the values in the right-hand panel of Figure 3.9, which were targeted in the model's calibration. In the baseline, the average follower-share is 8.6 percent, albeit with considerable heterogeneity across occupations. The red diamonds plot the follower-share when discounts are removed. The results are striking: following becomes much less likely in almost all occupations. The probability drops to less than one percent in the majority of cases. The largest change in the probability of following occurs for crop and animal producers (the third digit 6 occupation), where the share of followers drops from 15.3 percent to 1.8 percent. The counterfactual model predicts similar *percentage point* changes for Wood- and metal-plant operators (the first and second digit 8 occupations). The biggest *percent change* is predicted for religious professionals (the 16th digit 2 occupation), where following drops by 97.6 percent, from 6.0 percent to merely 0.1 percent.

The largest fall in 1-digit following occurs for professionals (digit

measure is highly correlated with educational requirements, making the inclusion of both difficult, absent any identifying restrictions.

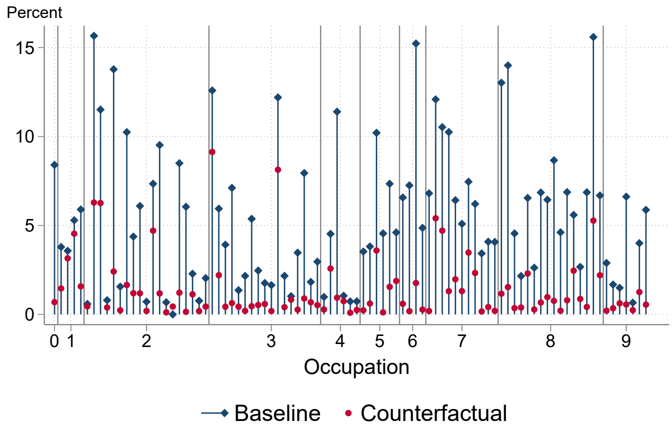

Figure 3.13: *Following in the Counterfactual Economy*

Note: This figure shows the fraction of fathers whose child follows them into the same occupations, for each occupation. The blue diamonds represent this fraction for the baseline model, the red circles report results for the counterfactual economy. On the x-axis, occupations are ordered according to their 3-digit code in the SSYK-96 classification system, the horizontal lines mark the borders of 1-digit occupational groups.

2) (see the gray squares in Figure 3.29 in Appendix 3.D). In the baseline model and the data, 42 percent of all sons born into this broad occupational category followed their father into the same category. After the removal of the discounts, this number drops to 27 percent. The decrease is similar for digit 6 occupations, agriculture and fishery workers, where the one-digit following probability falls from 16 percent to 3 percent.

Investigating the fall in occupational following across the father income distribution, Figure 3.14 shows that the effects are strongest in the first decile. Here, in the baseline economy, 9 percent of sons follow

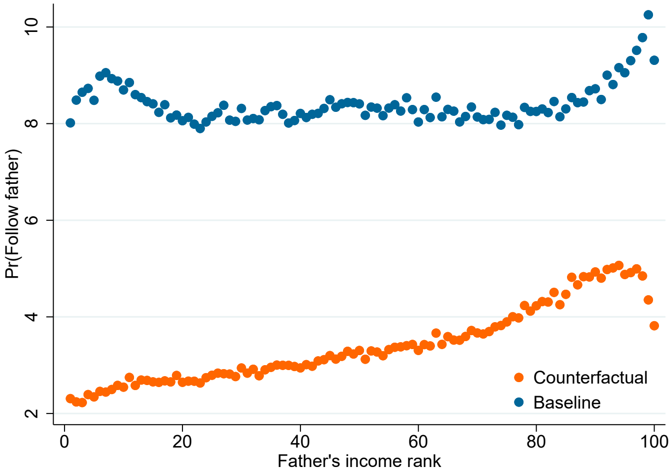

Figure 3.14: *Occupational Following by Father's Income Rank – Baseline and Counterfactual Economies*

Note: The figure shows the percentage of sons that follow into the same occupation as their fathers, across fathers' income bins, in the baseline and counterfactual economies, respectively.

their fathers, while only 3 percent do so in the counterfactual economy. Towards the top of the income distribution, even after the removal of the discounts, sons are disproportionally more likely to follow their fathers, compared to their peers with lower-income fathers.

From Figures 3.13, 3.14 and 3.29, we conclude that removing the discounts for followers strongly affects sons' occupational choices and pushes them away from their father's occupation. Next, we analyze how this affects intergenerational persistence in earnings.

3.7.2 Effects on Intergenerational Mobility

Given that sons are now less likely to follow their fathers into the same occupation, a natural question is whether removing the entry cost discounts increases earnings mobility in the economy. Figure 3.15a investigates this question by plotting the income rank-rank relationship between sons and fathers for the baseline and the counterfactual. The difference between the two is barely detectable, except at very low father income ranks. The correlation in income ranks between fathers and sons decreases from 0.291 to 0.245. The slope of the rank-rank graph is only slightly flatter due to this. The Gini coefficient decreases marginally, by about two percent, from 0.134 to 0.131.

Going beyond Figure 3.15a, which is informative about the changes in *ranked* earnings of sons, Figure 3.15b shows the percentage changes in sons' *real* earnings themselves, depending on the father's income rank. To compute the change in real earnings, we calculate each individual's nominal earnings in the baseline and counterfactual economies, respectively, and divide them by their respective prices indices.[36] Although the removal of parental influence on occupational choices of children is a major intervention to the baseline economy, the effect on earnings is small. The figure reveals that, in the counterfactual economy, real earnings increase most at the bottom of the distribution, where sons gain about two percent in terms of their prime-age earnings. Until percentile 60, real earnings changes are positive, but smaller; in the top decile they fall by more than one percent.

Beyond earnings changes, we can quantify how far sons move away in occupational space once discounts are removed. To this end, we first

[36] See Section 3.5.3

quantify the distance between each occupational pair in our sample using the Manhattan distance between their skill requirements according to the O*Net database, following Macaluso (2017).

Recall that each individual in our model has mass 1 which is potentially distributed across all occupations (due to the preference shocks). Thus, when moving from the baseline to the counterfactual economy, occupational changes do not occur discretely, i.e., from one occupation to another, but rather as a change in an individual's mass distribution across occupations. To quantify the distance between these distributions, we proceed in two steps. First, we take their difference, to determine how much mass is shifted. We assume that all *reductions* in mass allocated to occupations (when moving from baseline to counterfactual) are distributed randomly to those occupations which *gain* mass in the counterfactual economy. Then, in the second step, we determine how far, on average, the mass lost in each occupation travels to the new occupations. We take the average of these distances, within each sender occupation, across all receiver occupations, weighted by the share mass received in each occupation. This procedure quantifies, for each sender occupation, the average distance to receiver occupations. The final step is to average these distances across sender occupations, weighted by the share of total mass sent. This gives us, at the individual level, the average distance the shifted mass traverses when moving from baseline to counterfactual economy.

The left panel of Figure 3.16 shows the average occupational distance between the occupational choices in the baseline and the counterfactual economy. Distances are standardized across the population, implying that, in the first decile of the father income distribution, sons move roughly 0.2 standard deviations further away from their baseline occupational choices than the average individual in our sample.

Interestingly, occupational distances display a pronounced U-shape. Sons to fathers around the median income rank stay about 0.1 standard deviation closer to their baseline choices than does the average individual, but sons towards the top move much further from their baseline choices.

The measure of occupational distance reported in the left panel of Figure 3.16 combines the mass of the distribution which changes in the counterfactual economy (extensive margin) and the distance it travels from sending to receiving occupations (intenstive margin). The right panel of Figure 3.16 isolates the intensive margin, by holding the moving mass constant. This measure is largest at the low end of the father income distribution, where individuals move between 0.1 and 0.2 standard deviations further than the average individual, conditional on moving at all. At the top of the distribution, the picture is reversed, as individuals with fathers in the top percentile move 0.3 standard deviations less far away. Importantly, the extensive margin effect dominates the estimation of overall moving distance. Although sons to richer fathers stay closer to their original occupational choices conditional on moving, they are more likely to move overall.

In summary, Figure 3.16 suggests that, in response to the elimination of following discounts, (i) individuals with poor fathers move to occupations that are more different from their original occupations than individuals with rich fathers, and (ii) sons of rich fathers have a higher propensity to leave their baseline occupations when discounts are removed. The rich are more likely to move, but they stay close, conditional on moving.

3.7.3 Effects on Aggregate Output

Output, computed as aggregated real earnings, is 0.06% higher in the counterfactual economy compared to the baseline economy. An important component of this result is the general equilibrium nature of our model. In partial equilibrium, i.e., a hypothetical economy without discounts in which prices are not allowed to adjust, output actually decreases by 0.03%. The reason for the fall in aggregate earnings is that once discounts are removed, (i) blue-collar workers move into white-collar occupations, and (ii) occupations that pay high wages but require large entry costs, such as CEOs and medical doctors, see a sharp drop in entrants. The second effect dominates, which reduces aggregate earnings. According to the model, discounts serve two purposes in the baseline model: first, to compensate sons for relatively low earnings, as is the case in blue collar occupations, and second, to compensate sons for large entry costs which lower the utility of certain high-paying occupations. Consistent with these facts, the total amount of entry costs paid decreases in the partial equilibrium economy, relative to the baseline, by 7 percent.

The large influx into lower-paying white collar occupations, however, is not compatible with constant expenditure shares, as dictated by our model. Hence, in general equilibrium, the prices for the services provided by white-collar occupations decrease, with the notable exception of CEOs, lawyers and medical doctors, while those of blue-collar occupations increase. Total entry costs paid in the counterfactual economy are close to identical to the baseline economy, implying that the number of entrants into different occupations is very similar.

(a) Rank-Rank Relation – Baseline and Counterfactual Economies

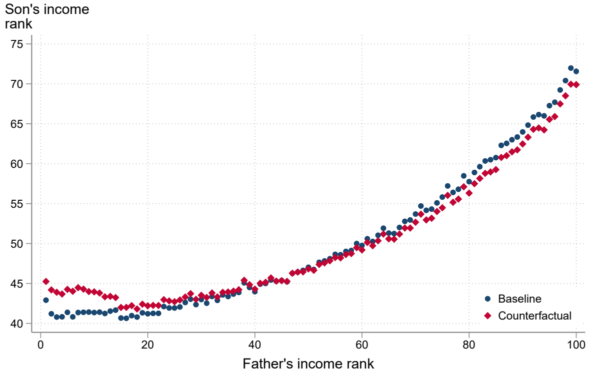

(b) Real Earnings – Counterfactual vs Baseline

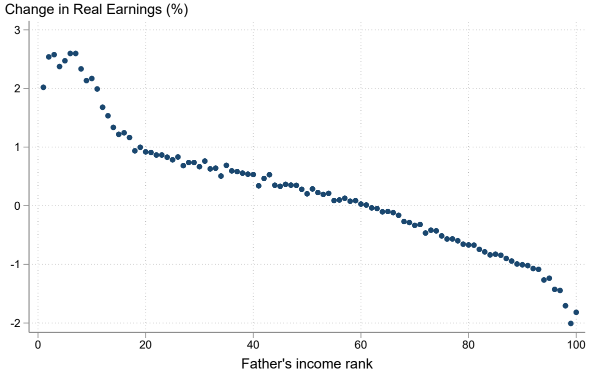

Figure 3.15: *Sons Earnings in Baseline and Counterfactual Economies*

Note: The figure shows sons earnings in baseline and counterfactual economies. Panel (a) plots the relationship between son's and their fathers' income ranks. Fathers are placed into 100 percentile bins. For each income bin for fathers, we calculate the average income rank of the sons, which is plotted on the y-axis. *Navy-colored dots* are based on results from the baseline model and the *red-colored diamonds* are based on the results from the counterfactual model. Panel (b) shows the average change in sons' real earnings, between the baseline model and the counterfactual, conditional on the income ranks of fathers. Fathers are placed into 100 percentile bins. For each income bin for fathers, we calculate the average earnings change for sons, which is plotted on the y-axis. Real earnings in the counterfactual economy include price effects.

(a) *Occupational distance moved* ***(b)*** *Normalized occupational distance moved*

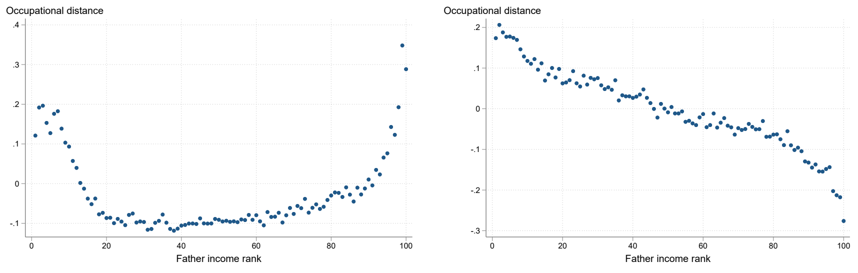

Figure 3.16: *Occupational distance between baseline and counterfactual*

Note: The *Left Panel* shows the average, standardized occupational distance travelled for an individual from the baseline economy to the counterfactual, by the income ranks of fathers. The *Right Panel* shows the average standardized occupational distance across the income ranks of fathers, holding the moving mass constant. For details, see the text.

3.8 Quasi-Experimental Evidence

The ideal experiment for answering the central question we ask in this paper, i.e., whether occupational following reflects misallocation of talent, would be one that level the playing field for children to enter all occupations, irrespective of the occupation of their parents. That is, a removal of any advantage of entering, or barriers of exiting, the occupation of their parent. We are not aware of any natural experiment of this kind. Therefore, to obtain such a counterfactual, we have estimated a structural model that allows us to perform this ideal counterfactual experiment, as presented in Section 3.7.

The results of the counterfactual experiment depend, however, on the model structure. Although our model matches important untargeted moments in the data, the model naturally requires assumptions about the structure of the labor market. Therefore, we present quasi-experimental evidence on the impact of parental occupation on occupational choice and income that is complementary to our model results. We are then able to obtain the same type of estimate using our structural model as the estimate relying on quasi-experimental variation.

3.8.1 Employment Decline in Father's Occupation

For a lack of an ideal natural experiment, we study the consequences of an employment decline in the fathers' occupations on the occupational choices and labor market outcomes of sons. We hypothesize that children whose father's occupation is in decline are less likely to follow into that occupation, e.g., due to reduced labor demand, lower wages, or weakening of father's network of contacts as the occupation shrinks. This, in turn, will serve as a first stage for an instrumental variable estimation of occupational following on labor earnings, where we use occupational decline as an instrument for a reduction in following.

For every son, we measure employment change in his fathers occupation as the change in the share of workers employed in the occupation out of total employment since fathers prime age. As we document in Appendix Figure 3.33, decline in employment in fathers' occupations is strongly correlated with advancement of labor-saving technologies in the occupation, measured either by the probability of occupations disappearing due to computerization (Frey and Osborne, 2017) or share of tasks automatized with robots (Webb, 2019). Since the intensity and root causes of occupational following may vary across occupations, and these are time-invariant measures, we proceed using a difference-in-differences strategy, exploiting the variation in employment change across fathers' occupations and across cohorts of sons. We estimate the following regression:

$$y_{int} = \alpha_n + \beta \Delta emp_{int} + \delta_t + \mathbf{X}_i' \gamma + \varepsilon_{int} \tag{3.14}$$

where y_{int} is the outcome of interest, e.g., the propensity of individual i to follow his father into occupation n, α_n are fathers' occupation fixed effects, Δemp_{nt} is change in employment in fathers occupation, δ_t are year-at-prime-age (i.e. bith cohort) fixed effects, and $\mathbf{X}_i$ is a vector of controls, including number of siblings and sibling order. The coefficient of interest is β, which measures the effect of employment change on the outcome of interest.

Figure 3.17 provides a graphical representation of the difference-in-differences regression estimate. First, it plots a binned scatter of the propensity to follow and the change in employment share in fa-

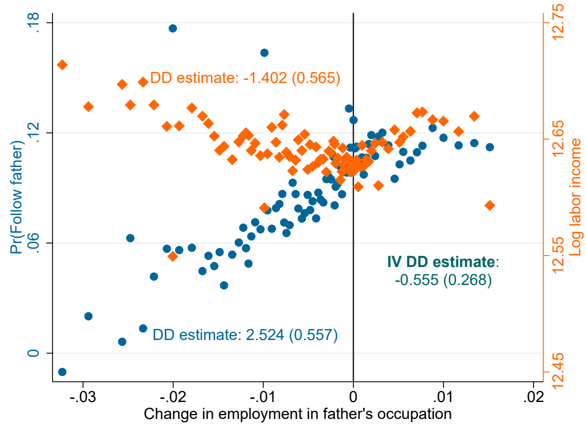

Figure 3.17: *Employment Decline, Occupational Following, and Labor Income*

Note: The figure plots the relationship between (i) the change in employment share in fathers' occupation since prime age and (ii) both the propensity of sons following into same occupation as their father (left axis) and labor earnings at prime age (right axis). The figure is a graphical representation of difference-in-differences regression (3.14) as it plots a binned scatter plot controlling for occupation and year-at-prime-age (cohort) fixed effects, as well as demographic controls including sibling indicator, and birth order dummies. *DD estimate* presents an estimate of β in regression (3.14). *IV-DD estimate* presents an estimate of a difference-in-differences regression as (3.14) but where log income is regressed on occupational following instrumented with occupational decline. Robust standard errors, clustered at father's occupation level, are in parentheses.

thers' occupation Δemp_{ot} controlling for father's occupation and cohort fixed effects, as well as demographic controls. In line with our hypothesis, there is a strong negative relation between occupational following and employment decline. The estimate of β is 2.5, implying that a 1 percentage point decline in employment in fathers' occupation as share of total employment leads to a reduction in occupational following by 2.5 percentage points. Second, Figure 3.17 also plots a binned scatter of log labor income and employment change in father's occupation. The estimate of β is -1.4, implying that a 1 percentage point decline in employment of fathers' occupation leads to about 1.4 percent increase in earnings. Under an exclusion restriction that a decline in fathers' occupation affects future earnings of sons only through occupational choice, occupational decline can be used as an instrumental variable for occupational following in a regression of earnings on following. The IV DD estimate is -0.55, meaning that sons who do not follow into their fathers occupation as a result of an employment decline in that occupation earn about 50 percent more than they otherwise would. This implies that those induced to switch an occupation away from that of their father move to occupations where they receive higher returns on their skills.

To evaluate this result, we divide sons into groups depending on their skills and their family background. Figure 3.18 presents the results. First, we divide sons into two groups depending on their skill-fit to their fathers occupation. More precisely, we split sons into those with a skill match to their father's occupation—measures by their predicted entry probability—above median and those that rank below median. For both groups the effect on the propensity to follow is negative and of similar magnitude. However, the earnings gain for sons from choosing another occupation than their fathers' are driven by sons whose skills are a relatively worse fit to that occupation. Second, we split sons in two groups by their father's income. Recall that the results in Section 3.7 documented that a decline in occupational following led to an earnings gains among sons of low-income fathers, but an earnings loss for sons of high-income fathers. Figure 3.18 documents that the impact of occupational decline on the propensity

Figure 3.18: *Effect of Employment Decline by Skills and Background*

(a) *Occupational Following* ***(b)*** *Labor income*

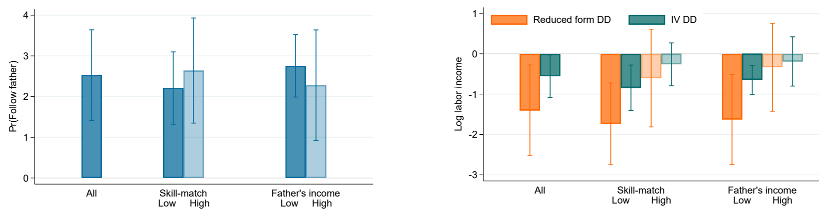

Note: This figure reports difference-in-differences regression estimates by individuals skills and family background. Panel (a) plots the estimated effect of employment decline on the propensity to follow into father' occupation, i.e. the first stage. Panel (b) plots the estimated effect of employment decline on labor earnings, i.e. the reduced form, and the IV estimate of following on labor earnings, where the propensity to follow is instrumented with the change in employment. We split the sample in two ways. First, by skill match to father's occupation, measured by the rank of predicted entry probabilities into the occupation. "High" and "Low" are indicators that refer to workers that have above or below median skill-fit. Second, we split workers by father's income, where "High" and "Low" are indicators that refer to sons whose fathers income is above or below median. Confidence intervals are based on robust standard errors, clustered at father's occupation level.

to follow is similar across both groups, but the effect on income is concentrated among sons of low income fathers. In sum, these results imply that occupational following among sons from poorer households reflect, at least to some extent, misallocation of talent.

3.8.2 Estimates Using the Structural Model

In our structural model, we can directly estimate changes in individuals' incomes and propensities to follow their fathers in response to changes in the following discounts. For an individual with a father in occupation n, these are the numerical derivatives of total income (which is affected through changes in the mass distribution across occupations) and following probability into occupation n (the mass allocated to occupation n), with respect to the discount of entering occupation n. The left panel of Figure 3.19 shows how the following

(a) *Derivative of following probability*

(b) *Structural model IV*

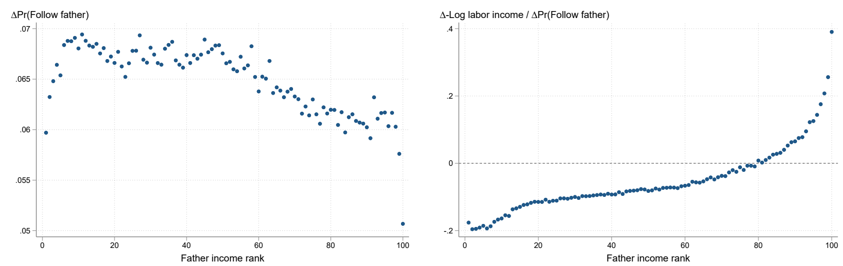

Figure 3.19: *Effects of discounts*

Note: The *Left panel* shows the change in following probabilities in response to a small change to following discounts. Results are averaged within father income ranks and scaled such that following discounts increase by the utility equivalent of 10,000 SEK. The *Right Panel* shows the ratio of the change in individual labor income and following probability, both in response to small changes in following discounts, averaged within father income bins. Fathers are placed into 100 percentile bins.

probability changes, in response to an increase in discounts equivalent to 10,000 SEK, for sons across the father income distribution. For sons of poorer fathers, this increases the following probability by around 7%; for sons with fathers in the top decile of the distribution, the effect is closer to 6%. Recall that average following probability in the baseline is about 8%, hence these results suggest that even moderate changes in following discounts can have large effects on sons propensities to choose their fathers occupations.

The right panel of Figure 3.19 constructs an estimate of how strongly labor income responds to changes in following probability, similarly to the IV estimate in the previous section. At the low end of the father income distribution, when people become more likely to follow their parents, due to increases in following discounts, their labor income decreases. Intuitively, they become more attached to occupations that rank towards the bottom of the earnings distribution. The pattern is reversed for sons of rich fathers: as they become more likely to follow, their labor income increases. These results are qualitatively in line with the empirical estimates from the previous section.

3.9 Conclusion

A child's economic success in adulthood is shaped by its parents in various ways. Inherited abilities and traits form an important foundation. However, environment and upbringing likely play important roles in the determination of the paths children take in their life. According to Old Norse wisdom, presented in Njáls saga, "*fostering makes the fourth part of a man*."

In this paper, we shed light on whether parental influence on the occupational choices of children leads to a misallocation of talent in the economy. First, we document that, across the whole range of occupations, children are likely to grow up to pursue the same occupations as their parents. Sons are about six times as likely to follow into their fathers' occupations as they would be if they were randomly assigned to occupations. Second, we estimate a structural general equilibrium Roy (1951) model that incorporates both heterogeneity in comparative advantage across workers due to their abilities and traits, as well as heterogeneous opportunities due to parental background. We find that in the absence of paternal influence on occupational choices through reduced entry cost, many sons choose to pursue a different occupation than their fathers. Still, social mobility and productivity are only marginally affected. These findings reflect the fact that while sons, with fathers in the different part of the income distribution, move to different occupations, they choose occupations which are on average similar to their original ones in earnings potential. Hence, the impact of parental influence on economic aggregates is small.

We also empirically estimate the effect of following on sons' earnings. Using a difference-in-differences approach that exploits cross-cohort variation in employment growth in father's occupation. Sons whose fathers' occupations are in decline are less likely to follow into that occupation, and, as a consequence of choosing another occupation, the sons that *do not* follow their fathers due to such decline *gain* in terms of prime-age income. This is driven by sons with skills that are the least aligned with their fathers occupation, implying that by following they do not exploit their full economic potential.

We emphasize that our results may not extend to all groups on the labor market. For example, women have, for a long time, faced unequal opportunities on the labor market, compared to men, due to unequal access to education, discrimination, and social norms. This is likely to have resulted in substantial misallocation of talent (Hsieh et al., 2019). Unfortunately, due to data limitation on measures of abilities and traits of women, which we obtain from tests associated with the military draft, the current paper is limited to the study of men's occupational choice. We hope to extend our focus in this direction in the future.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of labor economics*, volume 4, pages 1043–1171. Elsevier.

Adão, R. (2015). Worker heterogeneity, wage inequality, and international trade: Theory and evidence from brazil. *Working paper, MIT*.

Aghion, P., Akcigit, U., Hyytinen, A., and Toivanen, O. (2017). The social origins of inventors. Technical report, National Bureau of Economic Research.

Altmejd, A., Barrios-Fernández, A., Drlje, M., Goodman, J., Hurwitz, M., Kovac, D., Mulhern, C., Neilson, C., and Smith, J. (2021). O brother, where start thou? sibling spillovers on college and major choice in four countries. *The Quarterly Journal of Economics*, 136(3):1831–1886.

Autor, D. H. and Handel, M. J. (2013). Putting tasks to the test: Human capital, job tasks, and wages. *Journal of labor Economics*, 31(S1):S59–S96.

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

Becker, G. S. (1959). Union restrictions on entry. *The public stake in union power*, pages 209–24.

Becker, G. S. and Tomes, N. (1986). Human capital and the rise and fall of families. *Journal of labor economics*, 4(3, Part 2):S1–S39.

Bell, A., Chetty, R., Jaravel, X., Petkova, N., and Van Reenen, J. (2019). Who becomes an inventor in america? the importance of exposure to innovation. *The Quarterly Journal of Economics*, 134(2):647–713.

Bertrand, M. (2011). New perspectives on gender. In *Handbook of labor economics*, volume 4, pages 1543–1590. Elsevier.

Björklund, A. and Jäntti, M. (2012). How important is family background for labor-economic outcomes? *Labour Economics*, 19(4):465–474.

Black, S. E., Devereux, P. J., et al. (2011). Recent developments in intergenerational mobility. *Handbook of Labor Economics*, 4:1487–1541.

Black, S. E., Devereux, P. J., and Salvanes, K. G. (2005a). The more the merrier? the effect of family size and birth order on children's education. *The Quarterly Journal of Economics*, 120(2):669–700.

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

Black, S. E., Devereux, P. J., and Salvanes, K. G. (2009). Like father, like son? a note on the intergenerational transmission of iq scores. *Economics Letters*, 105(1):138–140.

Black, S. E., Grönqvist, E., and Öckert, B. (2018). Born to lead? the effect of birth order on noncognitive abilities. *Review of Economics and Statistics*, 100(2):274–286.

Black, S. E. and Strahan, P. E. (2001). The division of spoils: rent-sharing and discrimination in a regulated industry. *American Economic Review*, 91(4):814–831.

Blau, P. M. and Duncan, O. D. (1967). The american occupational structure.

Boar, C. and Lashkari, D. (2021). Occupational choice and the intergenerational mobility of welfare. Technical report, National Bureau of Economic Research.

Boserup, S. H., Kopczuk, W., and Kreiner, C. T. (2013). Intergenerational wealth mobility: Evidence from danish wealth records of three generations. *Univ. of Copenhagen mimeo*.

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

Carlsted, B. and Mårdberg, B. (1993). Construct validity of the swedish enlistment battery. *Scandinavian Journal of Psychology*, 34(4):353–362.

Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility i: Childhood exposure effects. *The Quarterly Journal of Economics*, 133(3):1107–1162.

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

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, 129(4):1553–1623.

Corak, M. and Heisz, A. (1999). The intergenerational earnings and income mobility of canadian men: Evidence from longitudinal income tax data. *Journal of Human Resources*, pages 504–533.

Dal Bó, E., Dal Bó, P., and Snyder, J. (2009). Political dynasties. *The Review of Economic Studies*, 76(1):115–142.

Deming, D. J. (2017). The growing importance of social skills in the labor market. *The Quarterly Journal of Economics*, 132(4):1593–1640.

Diamond, R. and Gaubert, C. (2021). Spatial sorting and inequality.

Doepke, M. and Zilibotti, F. (2008). Occupational choice and the spirit of capitalism. *The Quarterly Journal of Economics*, 123(2):747–793.

Edin, P.-A., Fredriksson, P., Nybom, M., and Öckert, B. (2022). The rising return to noncognitive skill. *American Economic Journal: Applied Economics*, 14(2):78–100.

Ermisch, J., Jantti, M., and Smeeding, T. M. (2012). *From parents to children: The intergenerational transmission of advantage*. Russell Sage Foundation.

Escriche, L. (2007). Persistence of occupational segregation: The role of the intergenerational transmission of preferences. *The Economic Journal*, 117(520):837–857.

Fredriksson, P., Hensvik, L., and Skans, O. N. (2018). Mismatch of talent: Evidence on match quality, entry wages, and job mobility. *American Economic Review*, 108(11):3303–38.

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

Ganzeboom, H. B., Treiman, D. J., and Ultee, W. C. (1991). Comparative intergenerational stratification research: Three generations and beyond. *Annual Review of sociology*, 17(1):277–302.

Gardberg, M., Heyman, F., Norbäck, P.-J., and Persson, L. (2020). Digitization-based automation and occupational dynamics. *Economics Letters*, 189:109032.

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

Gibbons, R. and Waldman, M. (2004). Task-specific human capital. *American Economic Review*, 94(2):203–207.

Goldin, C. (2014). A pollution theory of discrimination: male and female differences in occupations and earnings. In *Human capital in history: The American record*, pages 313–348. University of Chicago Press.

Goldin, C., Katz, L. F., and Kuziemko, I. (2006). The homecoming of american college women: The reversal of the college gender gap. *Journal of Economic perspectives*, 20(4):133–156.

Grönqvist, E., Öckert, B., and Vlachos, J. (2017). The intergenerational transmission of cognitive and noncognitive abilities. *Journal of Human Resources*, 52(4):887–918.

Gyllenram, A., Hellström, J., and Hanes, N. (2015). Förmåga att hantera stress och individers beslut att äga aktier. *Ekonomisk Debatt*, 43(1):7–15.

Härnqvist, K. (2000). Evaluation through follow-up. a longitudinal program for studying education and career development. *Seven Swedish longitudinal studies in behavioral science*, pages 76–114.

Hellerstein, J. K. and Morrill, M. S. (2011). Dads and daughters the changing impact of fathers on women's occupational choices. *Journal of Human Resources*, 46(2):333–372.

Hermo, S., Päällysaho, M., Seim, D., and Shapiro, J. M. (2022). Labor market returns and the evolution of cognitive skills: Theory and evidence. *The Quarterly Journal of Economics*.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The allocation of talent and us economic growth. *Econometrica*, 87(5):1439–1474.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1):35–78.

Kirkeboen, L. J., Leuven, E., and Mogstad, M. (2016). Field of study, earnings, and self-selection. *The Quarterly Journal of Economics*, 131(3):1057–1111.

Kramarz, F. and Skans, O. N. (2014). When strong ties are strong: Networks and youth labour market entry. *Review of Economic Studies*, 81(3):1164–1200.

Laband, D. N. and Lentz, B. F. (1983). Occupational inheritance in agriculture. *American Journal of Agricultural Economics*, 65(2):311–314.

Laband, D. N. and Lentz, B. F. (1985). *The roots of success: Why children follow in their parents' career footsteps.* Greenwood.

Laband, D. N. and Lentz, B. F. (1992). Self-recruitment in the legal profession. *Journal of Labor Economics*, 10(2):182–201.

Lentz, B. F. and Laband, D. N. (1989). Why so many children of doctors become doctors: Nepotism vs. human capital transfers. *Journal of Human Resources*, pages 396–413.

Lentz, B. F. and Laband, D. N. (1990). Entrepreneurial success and occupational inheritance among proprietors. *Canadian Journal of Economics*, pages 563–579.

Lindquist, M. J., Sol, J., and Van Praag, M. (2015). Why do entrepreneurial parents have entrepreneurial children? *Journal of Labor Economics*, 33(2):269–296.

Lo Bello, S. and Morchio, I. (2021). Like father, like son: Occupational choice, intergenerational persistence and misallocation. *Quantitative Economics.*

Long, J. and Ferrie, J. (2013). Intergenerational occupational mobility in great britain and the united states since 1850. *American Economic Review*, 103(4):1109–37.

Macaluso, C. (2017). Skill remoteness and post-layoff labor market outcomes.

Mayer, A. (2008). Education, self-selection, and intergenerational transmission of abilities. *Journal of Human Capital*, 2(1):106–128.

McFadden, D. (1974). The measurement of urban travel demand. *Journal of public economics*, 3(4):303–328.

Mocetti, S. (2016). Dynasties in professions and the role of rents and regulation: Evidence from italian pharmacies. *Journal of Public Economics*, 133:1–10.

Mood, C., Jonsson, J. O., and Bihagen, E. (2012). Socioeconomic persistence across generations: cognitive and noncognitive processes.

Murphy, K. M., Shleifer, A., and Vishny, R. W. (1991). The allocation of talent: Implications for growth. *The quarterly journal of economics*, 106(2):503–530.

Nakamura, E., Sigurdsson, J., and Steinsson, J. (2021). The gift of moving: Intergenerational consequences of a mobility shock. *The Review of Economic Studies*, Forthcoming.

Nicolaou, N., Shane, S., Cherkas, L., Hunkin, J., and Spector, T. D. (2008). Is the tendency to engage in entrepreneurship genetic? *Management Science*, 54(1):167–179.

Ohnsorge, F. and Trefler, D. (2007). Sorting it out: International trade with heterogeneous workers. *Journal of political Economy*, 115(5):868–892.

Rogoff, N. (1953). Recent trends in occupational mobility.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers*, 3(2):135–146.

Sattinger, M. (1975). Comparative advantage and the distributions of earnings and abilities. *Econometrica*, pages 455–468.

Sattinger, M. (1993). Assignment models of the distribution of earnings. *Journal of economic literature*, 31(2):831–880.

Svensson, A. (2011). Utvärdering genom uppföljning. longitudinell individforskning under ett halvsekel [evaluation through follow up. longitudinal individual research during half a century]. *Acta Universitatis Gotoburgensis*, 305.

Treiman, D. J. (1977). *Occupational prestige in comparative perspective*. New York: Academic Press.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Appendices

3.A Additional Material

3.A.1 Mapping Swedish Occupational Codes Over Time

The occupational codes in our dataset change over time. Before 1985, occupations are coded according to a three digit code named YK80; between 1985 and 1990, occupations are coded according to YK85, a five digit coding; and after that, occupations are coded according to SSYK96, a three digit coding. In order to facilitate our analysis, we elect to convert all codings into the most current one, SSYK96, at the three digit level.

We obtain a crosswalk between YK85 and SSYK96 from the Swedish statistical office (SCB). Conveniently, the former maps into the latter "m:1", i.e., multiple YK85 occupations map into the same SSYK96 occuption, but not vice versa.

The oldest occupational coding, YK80 also maps into SSYK96, but that mapping is "1:m", implying each of the older occupations maps into multiple recent ones. We tackle this problem by assigning each of the YK80 occupations exactly one SSYK96 counterpart, based on the highest overlap between the two. The tables describing crosswalks between the different occupational codings, produced by the Swedish statistical office, also indicate how many individuals assigned to occupation o in YK80 are assigned to each occupation P in SSYK96. In order to isolate a single SSYK96 occupation to which to assign each YK80 occupation, we pick the one to which most individuals are assigned, separately for men and women. We believe that this creates a credible crosswalk between the two codings. In almost 80 percent of all cases (for men), more than 70 percent of all individuals in a YK80 occupation are coded to one specific SSYK96 occupation and in 60 percent of all cases (for men), more than 90 percent of all individuals in a YK80 occupation are coded to one specific SSYK96 occupation.

3.A.2 Endogeneity of Skills

The data on skills used in this paper are based on measures at age 18. While these measures are intended to capture general skills, they may not reflect innate abilities. Instead, the skills and their measures may be influenced by the environment in various ways. Depending on how quantitatively important such endogeneity is, it could have important implications for our results. Importantly, if fathers invest in the skills of their sons that are most productive in their own occupation, and, in particular, if higher-income fathers engage more in such training than lower-income fathers, we may underestimate the true effect of parental occupation on intergenerational mobility. If skills are endogeneous in the way described above, we would expect that the relationship between the son's own skills and his father's skills and income would grow stronger over time.

To evaluate this concern, we leverage another source of data where individuals' skills are measured at younger ages. We use data on scores from tests administered as part of the *Evaluation Through Follow-up*, a large survey of Swedish families. These tests are taken when individuals are at age 12-13 (grade 6) and the data cover around 10 percent of the birth cohorts 1948, 1953, 1967, 1972, and 1977.[37] Härnqvist (2000) and Svensson (2011) provide details on the tests.

[37]The sample size of the survey, pooling across all cohorts, is roughly 20,000 inividuals.

Importantly, both data sources include tests for logical reasoning and vocabulary knowledge, which were unchanged across the cohorts.[38]

We restrict our sample to individuals from whom we have skills measured in both datasets. Restricting further to individuals for whom we also measure skills of their fathers reduces the sample substantially. We therefore report results both in terms of skills of fathers as well as father's income. We observe the number of questions that each person answered correctly on each test, both in the military enlistment and in the Evaluation Through Follow-up survey, out of a total of 40. We present results in terms of percentile (quintile) ranks when focusing on father's income (skills).

Figure 3.20 presents the intergenerational relationships between father's and son's skills, and between father's income and son's skills. Panel (a) plots the relationship between son's and father's logic-inductive ability, at ages 18 and 12/13. In line with inheritance of skills, there is a strong intergenerational correlation of skills. However, this pattern is remarkably similar at both younger and older ages, indicating limited effect of parental skills on their children's skills, above and beyond their initial inheritance. As explained above, the sample size is small where we have the triplet of skills measured at two ages for the sons and skills measured for their fathers. We therefore also present results where we relate skills of sons to income rank of fathers, which we can measure for almost all sons in the sample. As expected, there is a positive relationship between sons' skills and

[38]In the *Evaluation Through Follow-up* survey, the test on logical reasoning is to guess a number in a sequence of numbers, and the vocabulary knowledge test is to recognize antonyms (Svensson, 2011). In the military enlistment data, the logical reasoning test consisted of drawing correct conclusions based on statements that are made complex by distracting negations or conditional clauses and numerical operations, and the vocabulary knowledge test consisted of correctly identifying synonyms to a set of words (Carlsted and Mårdberg, 1993; Gyllenram et al., 2015).

Figure 3.20: *Comparison of Skills Measured at Age 18 and Age 12/13*

(a) *Logic-Inductive Ability by Fathers Skill Rank*

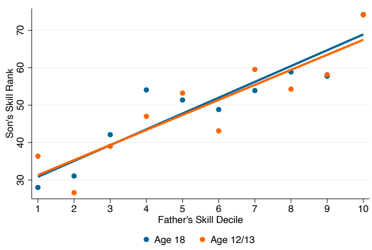

(b) *Logic-Inductive Ability by Fathers Income Rank*

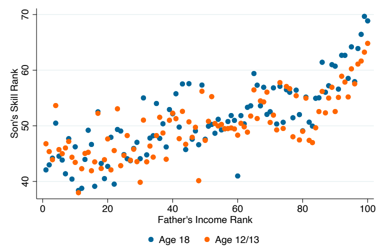

(c) *Verbal Comprehension by Fathers Skill Rank*

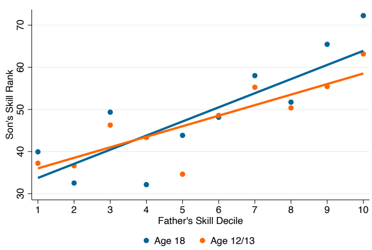

(d) *Verbal Comprehension by Fathers Income Rank*

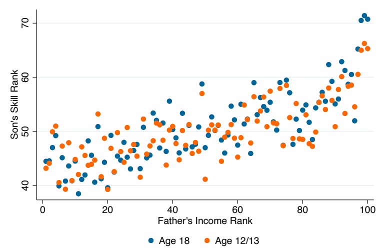

Note: This figure presents the intergenerational relationships between sons' and fathers' skills, and sons' skills and fathers' income rank. Skills are two cognitive skill measures: logic-inductive ability and verbal comprehension. Skills are measured at ages 12/13 (grade 6) and at age 18. The former is based on the *Evaluation Through Follow-up* while the latter is measured in tests administered as part of the military draft. The latter is our main measure used in this paper. Son's skills are measured as the percentile rank in their cohort. Father's skills are measured as a decile in the distribution of fathers within son's cohort, and father's income is measured as percentile rank in the distribution of fathers within son's cohort.

fathers income rank. As with father skills, this relation is almost the same when measured at ages 18 and 12/13. Panels (c) and (d) repeat the same exercise for the case of verbal comprehension, showing similar results.

We conclude from this exercise that we find limited evidence suggesting that skills of sons of high-skilled and high-income fathers change differently than that of lower-skilled and lower-income fathers over their early lives.

3.A.3 Occupational Following: Skills and Family Environment

The strong degree of occupational following documented in Section 3.3 may reflect three general phenomena. First, it may reflect selection on skills. This will be true if parents select into occupations based on skills that their children inherit (Roy, 1951), either genetically or by acquiring knowledge and experience through exposure to their parents' occupation while growing up (Becker and Tomes, 1986). Second, occupational following may reflect preferences. Parents may have a preference for their children following in their footsteps; children may have a preference for pursuing the same career as their parents, or, more generally, preferences over various aspects of occupations may be transmitted from parents to children (e.g., Escriche, 2007; Doepke and Zilibotti, 2008). Third, occupational inheritance may reflect unequal opportunities (Rogoff, 1953; Blau and Duncan, 1967). This can arise from unequal access to connections and social networks, access to information about possible careers, access to liquidity, or other aspects other than skills or preferences. Children of incumbents may then have an advantage when entering their parents' occupation or may face barriers to leaving it, due to financial or cultural frictions. While under the first explanation, children that follow their parents select into occupations in which they have a comparative advantage, the latter two explanations imply that occupational following may lead to misallocation of talent.

Figure 3.21 nonparametrically investigates the relationship between skill-fit, i.e., the entrance probability predicted by our machine learning algorithm, and a son's propensity to choose a given occupation. In all four panels of Figure 3.21, skill-fit is plotted on the x-axis. In order for the measure to be comparable across many occupations we generate percentile ranks of probabilities within occupations, such that those with the lowest entry probability have a rank of 1 but those with the highest have a rank of 100. Panel (a) documents that sons are more likely to enter their fathers' occupation the better their skills match to that occupation. While this may reflect selection on skills, an alternative explanation for this pattern is that there may be other factors than the skills we measure—but correlated to the skills—that influence sons' occupational choice. For example, as highlighted by Becker and Tomes (1986), parents may invest in their child's human capital, e.g., by training them to succeed in their own occupation. Similarly, fathers may, through contacts, be able to provide their sons with jobs in an occupation, regardless of their skills (Kramarz and Skans, 2014). To investigate this hypothesis, panel (b) of Figure 3.21 first restricts the sample to brothers (blue dots) and then partials out a father fixed effect (red diamonds). This leaves the relationship between the differences in brothers' skills and their differences in the propensity to follow. If the driver behind the pattern in panel (a) is family environment or training, we would expect the line of red diamonds to be flatter than blue dots: brothers should exhibit very similar propensities to follow. This is not the case. The introduction of fixed effects leaves the slope almost unchanged. We extend this analysis in Figure 3.21 by studying sons separately by birth order (panel c) and biological and adopted sons separately (panel d). While there is a strong relationship between skill match and following for all sons, first born sons are most likely to follow irrespective of skills, roughly 1 percentage points more likely than the second born and 2 percentage points more than the third born. This result speaks to prior studies documenting that earlier born children are more educated (Black et al., 2005a), have greater leadership skills, and are more willing to assume responsibility (Black et al., 2018), consistent with parents investing more in earlier than later born children. Lastly, in panel (d) we document that biological sons are 1.4 percent more likely to follow than adopted sons, but we still find a strong skill-gradient of following for both groups.

3.A.4 O*Net Skill Distance Robustness

As a validation exercise for our ideal occupation predictions, we construct measures of skill distance using them, which can be compared to measures of skill distance calculated using different data.

Macaluso (2017) estimates skill distance between two occupations using the O*Net database. Based on surveys, this dataset contains information on the average skillset of incumbents in each occupation, summarized as a 52-dimensional vector. She constructs the distance between occupations as the Manhattan-distance between the two skill vectors in each occupation pair.

First, following her approach, we construct the same measure for our dataset, after mapping the O*NET occupations into Swedish SSYK occuapations, as described in Section 3.A.4. Second, to construct the skill distance between two occupations i and j using **our** predictions, we do the following: Using our Random Forest algorithm, we ascertain, occupation-by-occupation, where an individual ranks, in terms of skill fit, *within an occupation*. Using this information, we calculate the Spearman correlation coefficient between the rankings of individuals for every occupation pair i and j in our dataset. If two occupations are more similar, we expect the fit-ranking of individuals to be more similar.

Figure 3.22, for medical doctors, shows a clear negative relationship between the skill distance estimated according to the O*NET data (Macaluso, 2017) on the x-axis and our measure of similarity on the y-axis. This gives us some confidence that our random forest algorithm is able to map skill sets into occupations faithfully.

Figure 3.23 plots the correlations between the different measures of skill distance across all occupations.[39] It is negative in almost all cases. The two approaches seem most consistent for the occupations including legislators and professionals, groups 1 and 2. Towards the blue collar occupations, while still negative, the two measures correlate less clearly.

Mapping International Occupational Codes into Swedish Codes

The O*NET database classifies occupations according to an SOC code. In order to map these into the Swedish SSYK96 system, we first map the SOC2010 code into an ISCO-08 code, which can then be mapped into SSYK2012, and finally into SSYK96.

The mapping between SOC2010 and the four-digit ISCO-08 classification is many-to-many. To calculate an ISCO-08 occupation's intensity in each of the different 52 different skills contained in the ONet database, we take the average of each of the skill measures across all

[39] Note that the O*NET database contains no information on military occupations.

SOC2010 occupations that map into it. For hypothetical ISCO occupation $\mathrm{I}-1$, we first find all SOC occupations that are linked to it, e.g., hypothetical occupations $\mathrm{S}-1$ and $\mathrm{S}-2$. To calculate the "oral comprehension" intensity of the $\mathrm{I}-1$ occupation, we take an average of the intensity in that skill across $\mathrm{S}-1$ and $\mathrm{S}-2$, weighted by the employment shares in $\mathrm{S}-1$ and $\mathrm{S}-2$.[40]. We proceed the same way for all other skills, e.g., "written comprehension" etc; and all other ISCO-08 occupations. Having done this, we obtain a dataframe containing the skill intensity for each of the ISCO-08 occupations, and all skills measured in the ONet database.

ISCO-08, in turn, maps into SSYK12 many-to-many. We use the same approach as before. First, to each SSYK12 occupation, we match all the ISCO-08 occupations that are linked to it. Then, we take the average over all the ISCO-08 occupations within each SSYK12 occupation, by skill. Thus, we obtain a dataframe containing the skill intensity for each of the SSYK12 occupations, and all skills measured in the ONET database.

From SSYK12 we proceed as in step one: merging SSYK12 to SSYK96 occupations and then obtaining average skill intensities for each skill-occupation pair by taking weighted averages, by SSYK12 occupation size.

3.A.5 Test Scores over Time

For our earnings predictions, we pool information about residualized incomes and test scores across several years (see Appendix 3.B for a description of how we pool data within 5-year periods). This approach would be problematic if the skill profiles needed to be well matched

[40]We obtain employment shares for all SOC occupations in 2014 from the BLS https://www.bls.gov/oes/tables.htm

to an occupation changed considerably over time.

In Figures 3.24 and 3.25, we graph the development of average cognitive and non-congnitive scores, respectively, for 1985 until 2013, in the four largest occupations in terms of employment. Although some trends are visible, e.g., that spatial ability has decreased over time among building finishers, the average scores are remarkably stable. That average test scores are stable within occupations across time is not unique to these four occupations. For each occupation, we calculate average scores for each of the eight respective tests, first for when we pool years 1985 and 1990, and then for years 2009-2013. We correlate the scores and find that the correlation coefficients are 0.96, 0.97, 0.89 and 0.93, respectively, for the cognitive tests, and 0.92, 0.84, 0.93 and 0.95, respectively, for the non-cognitive tests. This gives us confidence that our predictions are robust over time.

3.A.6 Prime Age Earnings by Occupation

Figure 3.26 documents the distribution of prime-age earnings across occupations. Not unexpectedly, average earnings are lowest for helpers in restaurants and highest for directors and chief executives, where the latter group earns roughly four times the earnings of the former. Across all occupations, weighted average earnings are 320 thousand SEK, with a standard deviation of 130 thousand SEK.

Figure 3.21: *Occupational Following – Skill Match and Family Background*

(a) *Occupational Choice – All Sons*

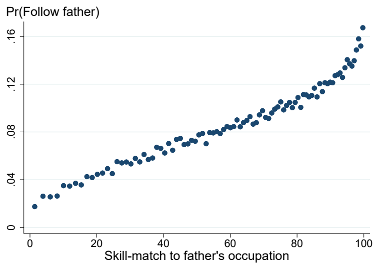

(b) *Brothers*

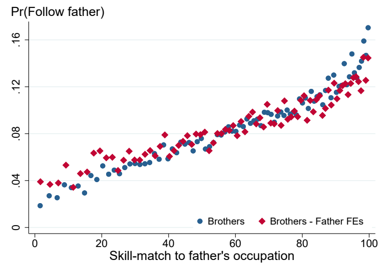

(c) *Brothers by Birth Order*

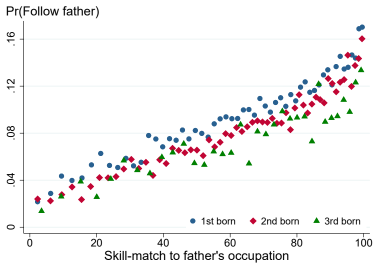

(d) *Biological and Adopted Sons*

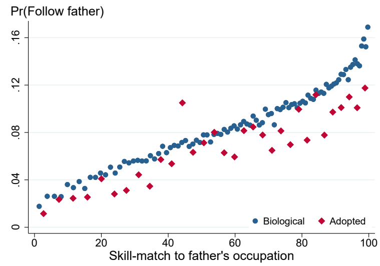

Note: This figure plots binned scatter plots of relationship between (i) the propensity to choose an occupation and (ii) the skill-match to that occupation, measured as the probability of entry predicted based on skills and presented in percentile ranks. All figures are based on regressions that partial out fixed effects for father's occupation. Panel (a) plots the relationship between skill-match and propensity to follow into father's occupation. Panel (b) plots the relationship between the propensity to follow and skill match for the sample of sons that have a brother in our sample, where dots show the raw relationship and diamonds show the relationship in differences across brothers, estimated using a regression including father fixed effect. Panel (c) plots the relationship between the propensity to follow and skill match by birth order for the sample of sons that have a brother in our sample. The group of "3rd born" sons includes third and later born sons. Panel (d) plots the relationship between the propensity to follow and skill match for biological and adopted sons.

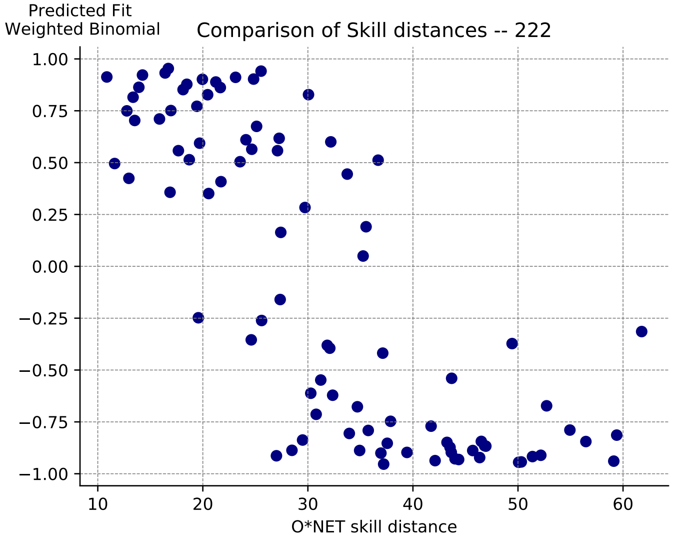

Figure 3.22: *Skill Distance and Occupation Similarity for Health Professions (except Nurses)*

Note: This figure shows the skill distance between two occupations, constructed according to Macaluso (2017), using O*NET data, on the x-axis and our measure of occupational similarity on the y-axis. The latter is the outcome of ranking all individuals according to their predicted entry probabilities (i.e., fit probabilities) in two different occupations and then calculating the Spearman correlation coefficient between the two rankings.

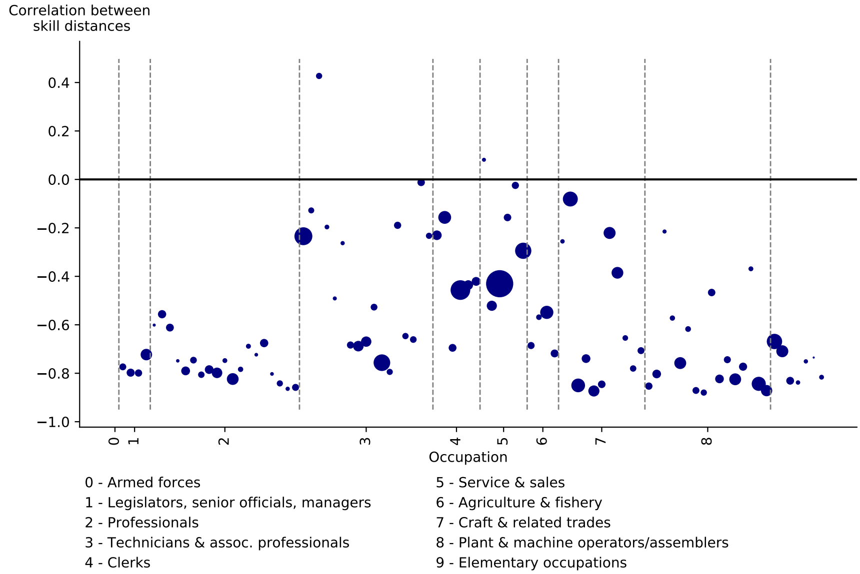

Figure 3.23: *Occupational Distance*

Note: This figure shows the correlation between two skill distance measures. The first is constructed according to Macaluso (2017), using O*NET data, the second is the outcome of ranking all individuals according to their predicted entry probabilities in two different occupations and then calculating the Spearman correlation coefficient between the two rankings. The y-axis in the figure shows the correlation between the two measures. On the x-axis, ccupations are ordered according to their 3-digit code in the SSYK-96 classification system, the vertical and horizontal lines mark the borders of 1-digit occupational groups.

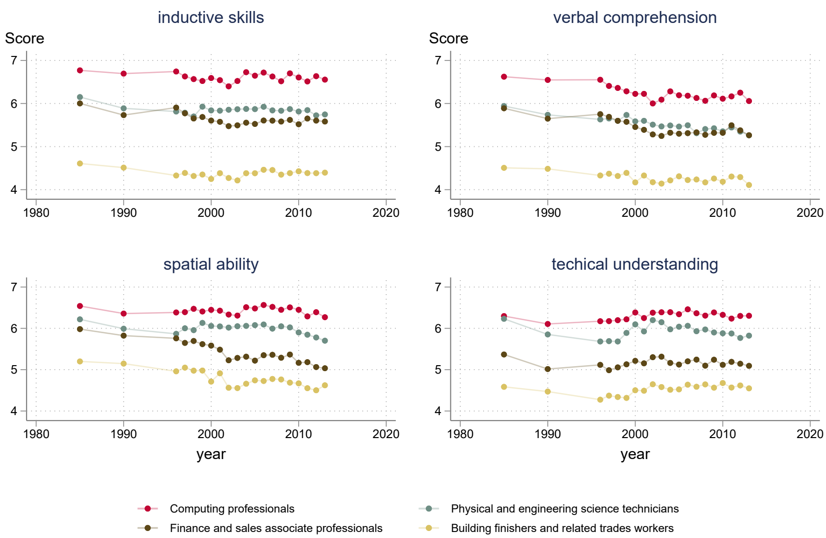

Figure 3.24: *Average Cognitive Test Scores of Incumbents over Time*

Note: This figure shows the average cognitive test scores of incumbents in four different occupations, over time. The *top left panel* shows average scores in inductive skills, the *top right panel* shows average scores in verbal comprehension, the *bottom left panel* shows average scores in spatial ability and the *bottom right* panel shows average scores in technical understanding. All figures feature incumbents in computing professionals (red), physical and engineering science technicians (green), finance and sales associate professionals (brown) and building finishers and related trades workers (yellow).

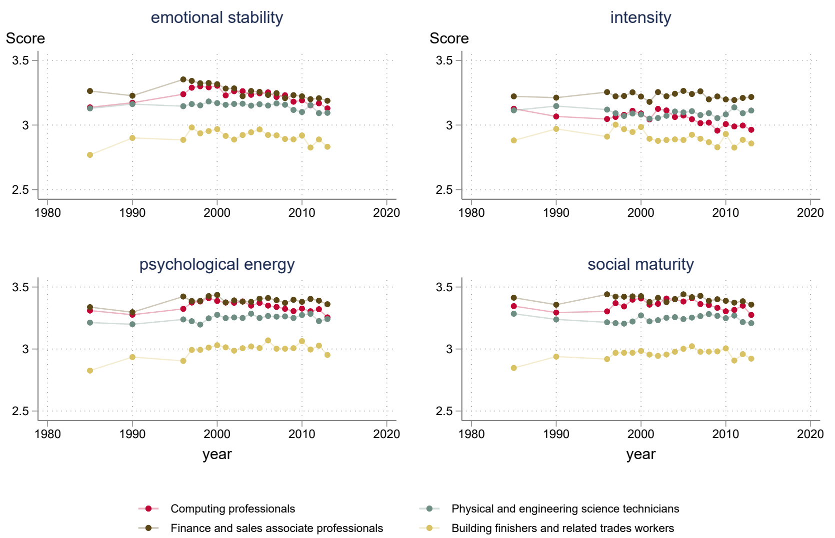

Figure 3.25: *Average Non-Cognitive Test Scores of Incumbents over Time*

Note: This figure shows the average non-cognitive test scores of incumbents in four different occupations, over time. The *top left panel* shows average scores in emotional stability, the *top right panel* shows average scores in intensity, the *bottom left panel* shows average scores in psychological energy and the *bottom right* panel shows average scores in social maturity. All figures feature incumbents in computing professionals (red), physical and engineering science technicians (green), finance and sales associate professionals (brown) and building finishers and related trades workers (yellow).

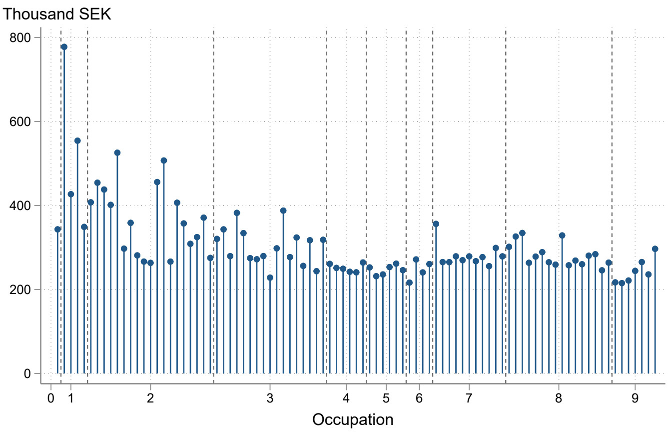

Figure 3.26: *Average Earnings – by Occupation*

Notes: This graph shows real average earnings, in 2013 SEK, by occupations in the period 1996-1999. On the x-axis, occupations are ordered according to their 3-digit code in the SSYK-96 classification system, the horizontal lines mark the borders of 1-digit occupational groups.

3.B Predictions of Probabilities of Occupation Entry and Income

In our analysis, we use skills to predict how well individuals fit into any given occupation they could choose from and what their potential earnings would be in a given occupation. To this end we use a random forest algorithm[41], which gives us flexibility to allow machine learning-driven selection of skills to be included in the prediction and their interaction.

3.B.1 Data Preparation

As the prediction is carried out sequentially by occupation (i.e. binomial as opposed to multinomial) we prepare two data sets for each occupation. The first is the training data and the second is the test data. In most cases the samples are the same, but in cases where we impose sample restrictions, such as building the prediction on the "best individuals" in each occupation, we impose those on the training data. The training data also has occupation-size weights which are used in the prediction. For each occupation, the dataset then has an indicator of whether individuals hold a given occupation or not (the outcome variable). We exclude sons who hold the same occupation as their father, i.e., followers from the prediction exercise.

3.B.2 Predicting Entry Probabilities

For each individual and each possible occupation, we predict the probability that the individual takes up that occupation based only on his skills. Training the algorithm on the incumbents in each occupation,

[41] We use the XGBoost package in **R**.

this, therefore, measures how well individuals fit into a given occupation. To account for the fact that occupations vary a lot in size, which will influence how accurately we can predict probabilities for small occupation, we use occupational-size weights in the model estimation.

The prediction process is a Random Forest estimation with cross validation (i.e. out-of-sample testing). The Random Forest algorithm is standard, where the number of splits are penalized if they do not yield a sufficient increase in prediction power. The cross-validation procedure works as follows:

1. The dataset X is split into n subsamples, $X_1, X_2 ... X_n$.

2. The XGBoost algorithm fits a boosted tree to a training dataset comprising $X_1, X_2, ..., X_{n-1}$, while the last subsample, X_n is held back as a validation (out-of-sample) dataset. The chosen evaluation metrics (RMSE) are calculated for both the training and validation dataset and retained.

3. One subsample in the training dataset is now swapped with the validation subsample, so the training dataset now comprises $X_1, X_2 ... X_{n-2}, X_n$, and the validation (out-of-sample) dataset is now X_{n-1}. Once again, the algorithm fits a boosted tree to the training data, calculates the evaluation metrics and so on.

4. This process repeats n times until every subsample has served both as a part of the training set and as a validation set.

5. Now, another boosted tree is added and the process outlined in steps 2-4 is repeated. This continues until the total number of boosted trees being fitted to the training data is equal to the number of rounds (i.e. the forest size).

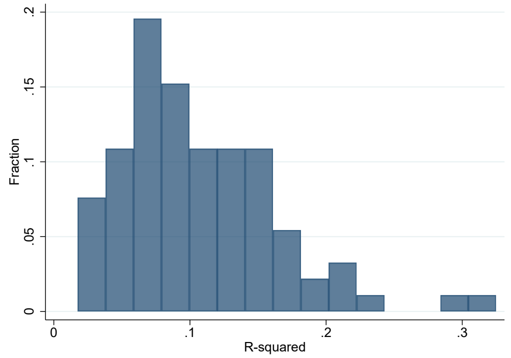

Figure 3.27: R^2 *across Occupation-Level Predictions*

Note: This figure plots the distribution of R^2 from random-forest predictions in each occupation. Prediction is based on the eight cognitive and non-cognitive skills of incumbents in each occupation. The sample period is 1985-2013.

6. There are now n calculated evaluation scores for each round for both the training sets and the validation sets. The prediction is then based on the round that best satisfies the evaluation metric (e.g. minimizes RMSE).

Based on this model, we then construct a predicted probability for all individuals in the given occupation. This procedure is then carried out for all occupations.

3.B.3 Predicting Income

The procedure for predicting income is analogous to the procedure for predicting probabilities, except for the fact that the prediction is linear as opposed to binary. The procedure is carried out separately for each occupation, yielding, for each individual, a predicted income in every occupation. The prediction is based on residualized income in logs. That is, we estimate the following regression:

$$\ln(earn_i) = \rho_o + \delta_c + \gamma_y + \varepsilon_i$$

where ρ_o, δ_c, and γ_y are, respectively occupation, birth cohort, and calendar year fixed effects. Then we use our machine learning approach to predict the earnings residuals across individuals and occupations. When translating the earnings predictions into SEK, we add fixed effects from the aforementioned regression. For comparability across the sample of individuals, we normalize earnings within each occupation by age and time, such that the reference age is 40 in a period. We split our sample into six periods, two per decade.

3.C Computation Appendix

3.C.1 Calibration of Baseline Economy

As described in Section 3.5.3, the baseline economy is calibrated to match data moments related to occupational choices. Costs and discounts are estimated jointly, as each of them affects all model moments. When we estimate the model, we do so in utility terms:

$$u(i,k) = \frac{Y(x(i),k)}{P} - b_k^f \tag{3.15}$$

where $Y(x(i),k)$ is the nominal income (and nominal expenditure) of individual i who works in occupation k, and P is the aggregate price index in the economy. b_k^f is the utility cost for entering occupation k, when individual i's father is in occupation f. See Equation (3.12) for more details.

We find initial guesses for our solution method as follows:

1) We consider entry costs only and target the share of sons in different occupations. The entry cost into military occupations is normalized to zero. Once we find an entry cost vector that yields shares that closely align with the corresponding data moments, we stop and store the vector as $m^{0,1}$.

2) Next, we target the shares of sons who choose the same occupational type (blue collar/white collar) as their fathers, taking $m^{0,1}$ as given. We iterate until we find that the model moments are close to their corresponding data moments. Call the resulting vector $d_1^{0,1}$. We normalize the discount for choosing a white-collar occupation to zero. This requires adjustments to the blue collar discounts and the entry cost vector, in order to

keep incentives the same. Label the adjusted vectors m^0 and d_1^0, respectively.

3) In the next step, we take take m^0 and d_1^0 as given and search for a vector of one-digit following discounts that brings the model close to the data. Once the model matches the data in this dimension, we store the resulting vector and call it d_2^0.

4) Last, we find a first guess for the set of follower discounts, holding all other discounts and costs fixed. We call this vector d_3^0. We normalize the follower discount into armed forces to zero.

Next, we iterate on all costs and discounts simultaneously, starting with the initial guesses obtained according to the above procedure, until the model moments match the data moments that we target. The estimated vectors are m, d_1, d_2, and d_3.

3.C.2 Counterfactual

In the counterfactual economy, we remove all discounts related to occupational following, and, following the use of the Cobb-Douglas aggregator for preferences, target the expenditure shares at their baseline values. To clear product markets, all prices $\{P_n\}_{n=1}^N$ adjust. For the baseline economy, we assumed that $P_n = 1\ \forall n$. As mentioned in Section 3.5.3, this normalization has no effect on relative predicted earnings across individuals within occupations, which is what matters for the results in the baseline economy. To find a new price vector $\{P_n^c\}_{n=1}^N$, given the entry costs m, estimated productivities $Z(x, n)$, and expenditure shares $\{\alpha_n\}_{n=1}^N$, we iterate on the price vector until the expenditure shares converge to the data values. As entry costs are measured in utils, we transform income to consumption utility by deflating nominal earnings by the price index $P^c = \prod_n \left(\frac{P_n^c}{\alpha_n}\right)^{\alpha_n}$, like in Equation (3.15).

3.D Additional Figures and Tables

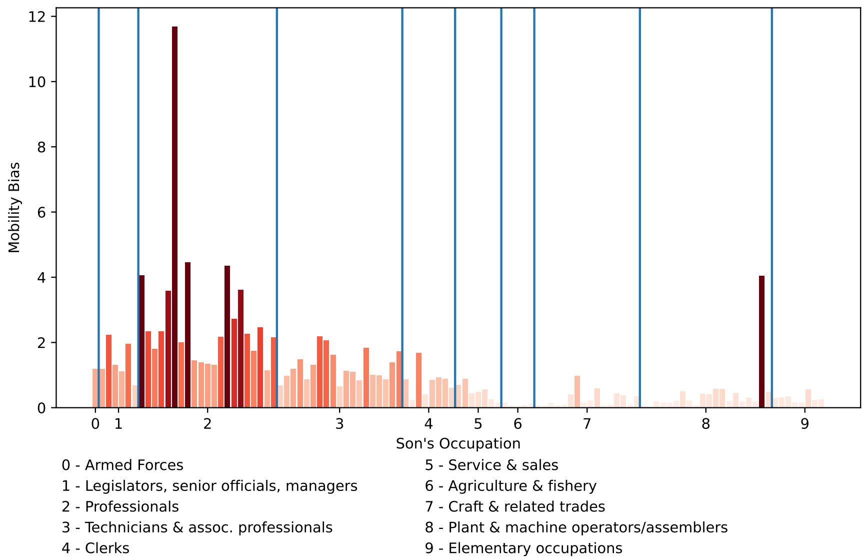

Figure 3.28: *Mobility Bias for Health Care Professionals – Occupational following*

Note: This figure shows the mobility biases on the y-axis for sons with fathers who are health professionals (except nursing). The values are equivalent to those reported in the associated row in Figure 3.2. On the x-axis, occupations are ordered according to their 3-digit code in the SSYK-96 classification system, the horizontal lines mark the borders of 1-digit occupational groups. For the computation of the mobility bias, see the text. The sample period is 1960-2013.

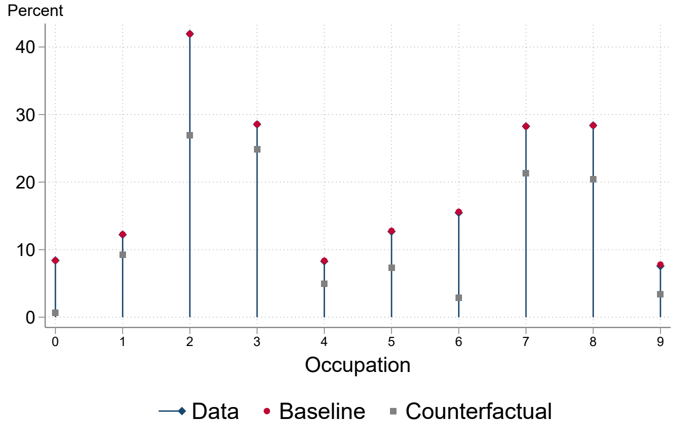

Figure 3.29: *Single Digit Occupational Following – Data, model, and counterfactual*

Note: This figure shows the fraction of fathers whose child follows them into the same broad occupational category, i.e., one-digit occupational classification. The blue diamonds represent this fraction for the pooled dataset, the red circles report the results for the baseline model and the gray squares report the results from our counterfactual exercise (see text). On the x-axis, occupations are ordered according to their 3-digit code in the SSYK-96 classification system, the horizontal lines mark the borders of 1-digit occupational groups. The sample period is 1985-2013.

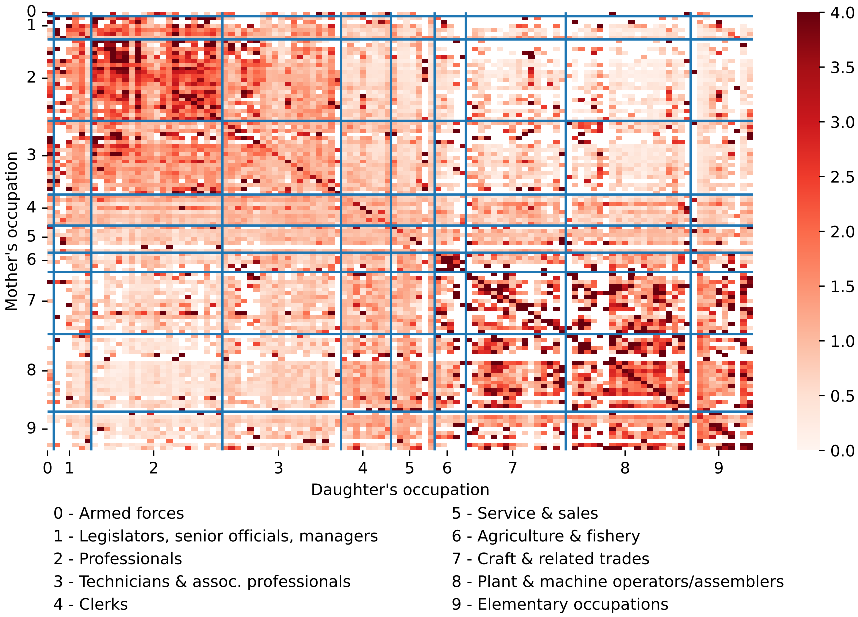

Figure 3.30: *Mobility Bias across Occupations – Mothers and Daughters*

Note: This figure shows the mobility bias estimates across different prime age occupations. The y-axis displays the mother's occupation, the x-axis displays the daughter's occupation. Occupations are ordered according to their 3-digit code in the SSYK-96 classification system, the vertical and horizontal lines mark the borders of 1-digit occupational groups. For the definition of the mobility bias, see the text. The sample period is 1960-2013.

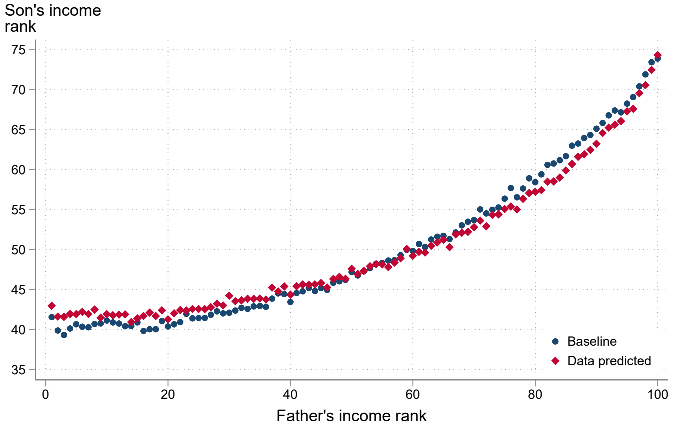

Figure 3.31: *Rank Rank Relation – Baseline and Data*

Note: The figure shows the relationship between son's and their fathers' income ranks. Fathers are placed into 100 percentile bins. For each income bin, we calculate the average income rank of the sons, which is then plotted on the y-axis. *Navy-colored dots* are based on results from the baseline model and the *red-colored diamonds* are based on empirical data.

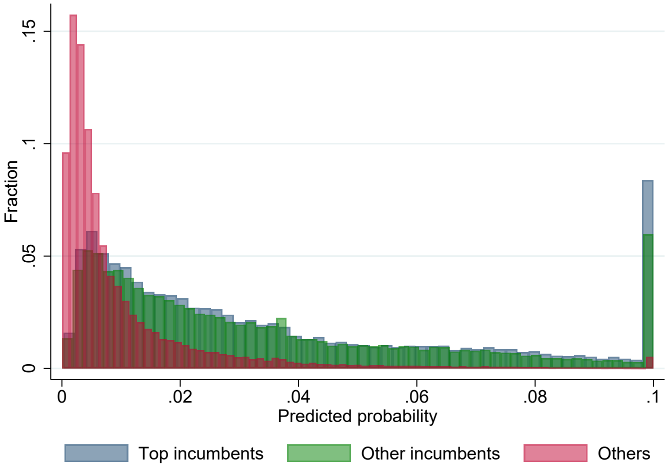

Figure 3.32: *Predicted Probability of Occupation Entry*

Note: The figure shows predicted probability of entry into occupations. The figure separated three groups: "Top incumbents" which are incumbents in the occupation in the top quintile of the earnings distribution and those used for training the machine-learning algorithm, "Other incumbents" which includes all other incumbents in the occupation, and "Others" which are workers in other occupations. The figure is winterized from above at 10 percent probability of entry.

Figure 3.33: *Occupational Decline: Automation and Robotization*

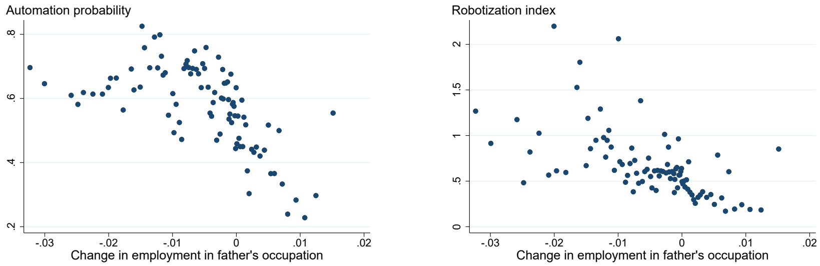

(a) *Automation* ***(b)*** *Robotization*

Note: The figure plots a binned scatter of the correlation between (i) a change in employment share in fathers' occupation for all sons in our sample and (ii) two measures of labor-saving technological change. In panel (a) we plot occupation-specific automation probabilities based on Frey and Osborne (2017). This measure is based on analysis of 702 US occupation and measures probability in 2010 that an occupation will disappear within 10-20 years due to computerization. Using this measure, Gardberg et al. (2020) also document a decline in employment share since the 1990s in occupations more exposed to risk of automation. In panel (b) we plot occupation-specific measure of exposure to automation measured by tasks that can be performed by industrial robots, as measured by Webb (2019).

Table 3.1: *List of Occupations: SSYK-96 Codes and their Descriptions*

SSYK96 code	Description
011	Armed forces
121	Directors and chief executives
122	Production and operations managers
123	Other specialist managers
131	Managers of small enterprises
211	Physicists, chemists and related professionals
213	Computing professionals
214	Architects, engineers and related professionals
221	Life science professionals
222	Health professionals (except nursing)
223	Nursing and midwifery professionals
231	College, university and higher education teaching professionals
232	Secondary education teaching professionals
233	Primary education teaching professionals
235	Other teaching professionals
241	Business professionals
242	Legal professionals
243	Archivists, librarians and related information professionals
244	Social science and linguistic professionals (except social work professionals)
245	Writers and creative or performing artists
246	Religious professionals
247	Public service administrative professionals
248	Administrative professionals of special-interest organisations
249	Psychologists, social work and related professionals
311	Physical and engineering science technicians
312	Computer associate professionals
313	Optical and electronic equipment operators
314	Ship and aircraft controllers and technicians
315	Safety and quality inspectors
321	Agronomy and forestry technicians
322	Health associate professionals (except nursing)
323	Nursing associate professionals
331	Pre-primary education teaching associate professionals
332	Other teaching associate professionals
341	Finance and sales associate professionals
342	Business services agents and trade brokers
343	Administrative associate professionals
344	Customs, tax and related government associate professionals
345	Police officers and detectives
346	Social work associate professionals
347	Artistic, entertainment and sports associate professionals
412	Numerical clerks
413	Stores and transport clerks
415	Mail carriers and sorting clerks
419	Other office clerks
421	Cashiers, tellers and related clerks
422	Client information clerks
511	Travel attendants and related workers
512	Housekeeping and restaurant services workers
513	Personal care and related workers
514	Other personal services workers
515	Protective services workers
522	Shop and stall salespersons and demonstrators
611	Market gardeners and crop growers
612	Animal producers and related workers
613	Crop and animal producers
614	Forestry and related workers
711	Miners, shotfirers, stone cutters and carvers
712	Building frame and related trades workers
713	Building finishers and related trades workers
714	Painters, building structure cleaners and related trades workers
721	Metal moulders, welders, sheet-metal workers, structural-metal preparers and related trades workers
722	Blacksmiths, tool-makers and related trades workers
723	Machinery mechanics and fitters
724	Electrical and electronic equipment mechanics and fitters
731	Precision workers in metal and related materials
734	Craft printing and related trades workers
741	Food processing and related trades workers
812	Metal-processing-plant operators
814	Wood-processing- and paper-making-plant operators
815	Chemical-processing-plant operators
816	Power-production and related plant operators
821	Metal- and mineral-products machine operators
822	Chemical-products machine operators
823	Rubber- and plastic-products machine operators
824	Wood-products machine operators
825	Printing- binding- and paper-products machine operators
826	Textile-, fur-, and leather-products machine operators
827	Food and related products machine operators
828	Assemblers
829	Other machine operators and assemblers
831	Locomotive-engine drivers and related workers
832	Motor-vehicle drivers
833	Agricultural and other mobile-plant operators
912	Helpers and cleaners
913	Helpers in restaurants
914	Doorkeepers, newspaper and package deliverers and related workers
915	Garbage collectors and related workers
919	Other sales and services elementary occupations
932	Manufacturing labourers
933	Transport labourers and freight handlers

Chapter 4

Monetary Policy and Liquidity Constraints: Evidence from the Euro Area*

*This chapter was jointly written with José-Elías Gallegos, John Kramer, and Ricardo Lima. We are grateful to Tobias Broer, Per Krusell, Kurt Mitman and Kathrin Schlafman for their advice and support. Further, we would like to thank our editor at the American Economic Journal: Macroeconomics, Giorgio Primiceri, four anonymous referees, Adrien Auclert, Martín Beraja, Mitch Downey, Xavier Gabaix, Alessandro Galesi, Pierre-Olivier Gourinchas, John Hassler, Greg Kaplan, Peter Karadi, Karin Kinnerud, Kasper Kragh-Sørensen, Jesper Lindé, Benjamin Moll, Haroon Mumtaz, Morten Ravn, Ricardo Reis, Giovanni Ricco, Federica Romei, Maria Sandström, David Schönholtzer, Jósef Sigurdsson, Xueping Sun, Javier Vallés, Has van Vlokhoven, Iván Werning and seminar participants at the IIES, Queen Mary University of London, Stanford University, Stockholm University, Stockholm School of Economics, Università Bocconi and Yale University for useful feedback and comments. Björn Hagströmer generously aided us in accessing the swap-rate data. José-Elías Gallegos and John Kramer greatfully acknowledge funding from the Tom Hedelius foundation. This paper uses data from the Eurosystem Household Finance and Consumption Survey. The results published and the related observations and analysis may not correspond to results or analysis of the data producers.

4.1 Introduction

In 2016, 30 percent of households in Germany reported that they could not meet an unexpected, immediate financial expense of 985 euros. At the same time, 40 percent of Italian households reported that they were unable to meet an unexpected expense of 800 euros.[1] Figures like these suggest that a significant portion of households hold little liquid assets, which potentially makes them vulnerable to unexpected shocks to the economy. Especially in monetary economics, these households have received special attention recently.

While theoretical research has shown that heterogeneous agent models which include constrained agents can have different policy implications than their representative agent counterparts, empirical evidence on how heterogeneity matters for the transmission from monetary policy to output is scant.[2] In this paper we provide such evidence, showing that a higher share of liquidity constrained households in a country is associated with a stronger output response to a monetary policy surprise.

We focus on the euro area, where member countries have been exposed to the common monetary policy conducted by the European Central Bank (ECB) since the introduction of their shared currency. However, because of long-standing country idiosyncrasies and slow convergence, they still differ along many dimensions, including the share of liquidity constrained households, as we show. Since we choose

[1]According to the European Union Survey of Income and Living conditions. The monetary values represent the country-specific at-risk-of-poverty threshold, defined as 60 % of the national median equivalized disposable income after social transfers.

[2]See e.g., Bilbiie (2008) for an early theoretical contribution in a two agent setting or Auclert (2019) and Hagedorn et al. (2019) for a setting with fully heterogeneous agents.

this "bird's eye view", we can conduct standard monetary policy analysis, while taking account of wealth and income heterogeneity and its influence on output responses.

First, we estimate output impulse response functions (IRFs) at monthly frequencies for each country to the same monetary policy shocks, relying on the Local Projection (LP) approach pioneered by Jordà (2005).[3] Because of endogeneity concerns between policy rate changes and output responses, we augment the LP estimation with an instrumental variable (IV) framework (Stock and Watson, 2018).[4] We use high-frequency movements in Overnight Indexed Swap (OIS) rates in a 45 minute time window around ECB policy announcements as an instrument for monetary policy surprises. Because OIS are forward looking interest rate derivatives, large rate movements during the window imply that the ECB's announcement was not in line with market expectations. The identifying assumption is that this measure is uncorrelated with other shocks to output.

In the second part of the paper, we incorporate the income and asset dimensions by relating the IRFs to the share of liquidity constrained households in each country. The idea is that a higher fraction of households less able to smooth the income fluctuations caused by monetary policy shocks may lead to a stronger aggregate output response in a country. While it is not possible to measure the fraction

[3]Mandler et al. (2016) investigate a similar question using a Bayesian VAR for the four largest economies in the euro area: Germany, Italy, Spain and France. Altavilla et al. (2016) investigate heterogeneous effects of Outright Monetary Transactions (OMT) on the same countries, similarly employing a VAR framework.

[4]As a robustness check to our main empirical framework, we construct an instrumented Global VAR (GVAR) based on Georgiadis (2015) and Burriel and Galesi (2018). We build a more structural –and restricted– setting than the LPIV, more similar to the widespread VAR estimation in the literature, identifying monetary responses in a GVAR setting using exogenous instruments. To our knowledge, we are the first to estimate such an instrumented GVAR. We find similar results.

directly, we approximate it by classifying households in the Household Finance and Consumption Survey (European Central Bank, 2014, HFCS) as Hand-to-Mouth (HtM) or non-HtM according to a measure proposed by Kaplan et al. (2014). They show that this measure is strongly correlated with estimates of marginal propensities to consume (MPC). Since the HFCS can only provide data on recent years, we complement it with data from the European Union Survey on Income and Living Conditions (EU-SILC), which has been conducted since 2005. In it, participating households are asked whether they could finance an unexpected financial expense, from which we infer whether they are financially constrained. Both surveys point to large variation across countries in the share of constrained consumers and the pattern is broadly consistent over time.

Our first finding is that, in line with previous literature, output responses to a common European monetary policy surprise are not homogeneous across countries. There is significant heterogeneity in terms of cumulative impact and peak values. Secondly, all of our measures of the fraction of liquidity constrained households are significantly correlated with the strength of the IRFs. On average, countries with higher fractions of liquidity constrained households exhibit stronger cumulative output responses and bigger peak responses to an unexpected interest rate change. For the measure constructed according to Kaplan et al. (2014), we show that the results are driven by the "wealthy HtM", i.e., households with low levels of liquid wealth, but positive and possibly large levels of illiquid wealth. In addition, we calculate aggregate output IRFs for a constrained and a less-constrained group of countries. The two responses are significantly different at most horizons, with the more constrained countries reacting more strongly to the common shock.

The results we present are important for several reasons. First, our findings suggest that heterogeneity in the composition of household balance sheets across countries affects the transmission of monetary policy to their economies. The finding that a higher share of low-liquidity households amplifies the output response to an unexpected interest rate change can guide future theoretical and quantitative work on monetary policy in a Heterogeneous Agent New Keynesian framework. Understanding the reasons for the differences we uncover is crucial in order to calibrate future policies. Second, we show that LP methods can be used to estimate the impact of monetary policy for countries within a currency union. Lastly, our results are robust across different specifications of liquidity constraints, corroborating the measure put forth by Kaplan et al. (2014).

Our research is related to several strands of literature. There is a large body of research which performs cross-country monetary policy analysis. An early example is Gerlach and Smets (1995) who perform a Structural VAR analysis of the G-7 countries and find that responses to country-specific monetary policy shocks are similar. Mandler et al. (2016), using a large Bayesian VAR, show that output in Spain is less responsive to monetary policy, compared to Germany, France and Italy, while prices in Germany respond most within this set of countries. Few papers estimate IRFs for multiple countries and try to investigate the transmission mechanism of monetary policy by relating their findings to country characteristics. Two recent examples, both of which use a Global VAR (GVAR) method, are Georgiadis (2015) and Burriel and Galesi (2018). Both papers find heterogeneous responses of real GDP across countries and explain some of the variation with wage rigidities and the fragility of the banking sector. Calza et al. (2013) provide evidence that in countries where the use of flexible mortgage rates is more prevalent, responses to monetary policy shocks are stronger and Corsetti et al. (2018) find that the responses of output and private consumption are larger in countries where home ownership rates are higher. We try to account for previous findings by conducting several robustness checks.

To our knowledge, we are the first to use OIS rates as an instrument to identify a cross-country LP estimation in the euro area. Kuttner (2001), Nakamura and Steinsson (2018) and Gertler and Karadi (2015) use high-frequency movements in Federal Funds futures rates in a short window around the Federal Reserve's policy announcements to identify monetary policy surprises in the U.S. In the European context, there are no financial instruments equivalent to Fed funds futures which has led researchers to employ high-frequency movements in OIS rates instead. Ampudia and Van den Heuvel (2018) and Jarociński and Karadi (2020) construct monetary policy shocks from movements in these derivatives.

The empirical results in this paper tie in with the results from theoretical two-agent New Keynesian (TANK) models such as those in Bilbiie (2008), Galí et al. (2007) and Bilbiie (2020), as well as richer models by Gornemann et al. (2016), Werning (2015), Auclert (2019) and Hagedorn et al. (2019). As laid out by Bilbiie (2019), a result these models have in common is that whether aggregate shocks have bigger or smaller effects on aggregate consumption, compared to the representative agent framework, is ambiguous. In a model that combines the tractability of TANK models with the most important elements of heterogeneneous agent models, Bilbiie (2019) shows that the output response to shocks is amplified if the income elasticity of constrained agents with respect to aggregate income is larger than one. He refers to this case as cyclical income inequality; a channel which is strengthened if a larger fraction of agents is constrained.[5] This is in line with our empirical findings, which can guide future modeling efforts aimed at understanding the interaction between aggregate and distributional outcomes in response to shocks.

Lastly, our findings imply that it is important to separately treat liquid and illiquid assets when describing the wealth distribution of an economy. This is in support of the view that wealthy households can have high marginal propensities to consume, as pointed out by Kaplan et al. (2014), Kaplan and Violante (2014) and Kaplan et al. (2018).[6]

The paper proceeds as follows. In section 4.2, we describe our identification strategy, how we estimate country-specific local projections and present the resulting IRFs. Section 4.3 discusses how we construct the proxies for the fraction of liquidity constrained households across countries. Section 4.4 relates the IRFs to our measures of the fraction of liquidity constrained households across countries. Section 4.5 concludes.

[5]In models that focus on the cyclicality of income risk (e.g., Werning, 2015), amplification of aggregate shocks is caused by an increase in the probability of becoming constraint for the unconstrained, which leads the latter to save more and consume less. Our empirical analysis, however, focuses mainly on the level of the HtM shares, as opposed to their changes, and is therefore more closely related to Bilbiie (2019).

[6]Using data from Norwegian lottery winners, Fagereng et al. (2020) find that households at the highest liquidity quartile have a significantly lower MPC than households at the lowest liquidity quartile.

4.2 Effects of Monetary Policy Shocks

4.2.1 Identifying monetary policy shocks

In order to estimate the effects of monetary policy on a variable of interest, we need to identify unexpected deviations from an interest rate rule. To identify these in the United States, researchers have used high frequency movements in Federal Funds futures in a narrow time window around announcements by the Federal Reserve (Kuttner, 2001; Nakamura and Steinsson, 2018). More recently, Ampudia and Van den Heuvel (2018) and Jarociński and Karadi (2020) apply the approach to European data using Overnight Indexed Swap (OIS) rate movements around ECB announcements. These derivatives are traded over-the-counter between two parties exchanging a fixed interest rate for the floating Eonia overnight interest rate, both on a notional principal, for a pre-specified amount of time. Since the principal is not exchanged at any time and the contracts are highly collateralized, there is only minimal counterparty credit risk. When the contract ends, the difference between (i) the fixed interest accrued on the principal and (ii) the interest accrued on the principal by investing it at the overnight interest rate every day is calculated and the contract is cash settled.[7]

We follow the literature and use changes in Eonia OIS during a short time window around the ECB's monetary policy announcement and the subsequent press conference as our instrument (Jarociński and Karadi, 2020).[8] On days when the Governing Council of the ECB de-

[7]For a detailed discussion of similarities between federal funds futures and overnight indexed swaps, see Lloyd (2018).

[8]We obtain time series data on OIS rates at the minute frequency from Datascope (2018), using the #RIC *EUREON3M=* and *EUREON1Y=*. The time format is GMT/UTC. For more information, see Appendix 4.E.

cides the policy rate for the euro area, the decision is communicated to the public via a press statement at 13:45 CET and motivated during a press conference chaired by the president and vice-president at 14:30 CET. We construct a time series encompassing all such monetary policy announcements by the ECB, starting in December 1999.[9] Figure 4.1 displays the OIS rate on July 5, 2012. The first window starts 15 minutes before the press release and ends 30 minutes after. The second window starts 15 minutes before the beginning of the press conference and ends 30 minutes after. To construct our instrument, on each announcement date, we calculate the change in the average OIS rates of the pre- and post-windows for both the press statement and the press conference and then sum the two.

The OIS rate can be viewed as an indicator for expectations about future overnight interest rates in the European interbank market. Hence, a significant change in the OIS rates shortly after an ECB monetary policy announcement implies that the content of the announcement was at least partly unexpected. The identifying assumption is that there is no other information released during the time window that is systematically related to the policy decision, and that the market has access to the same information about economic fundamentals as the ECB. As pointed out by Jarociński and Karadi (2020), many of the Bank of England's announcement dates coincide with announcement dates of the ECB, with policy statements released at 13:00 CET and 13:45 CET, respectively. This makes the high-frequency approach especially important in our setting. The use of instruments measured at the daily frequency would confound the effect of the former and the latter.

[9] To construct monetary shocks starting from January 2000, we start collecting movements in OIS rates from December 1999, due to the way we construct our instrument (see below).

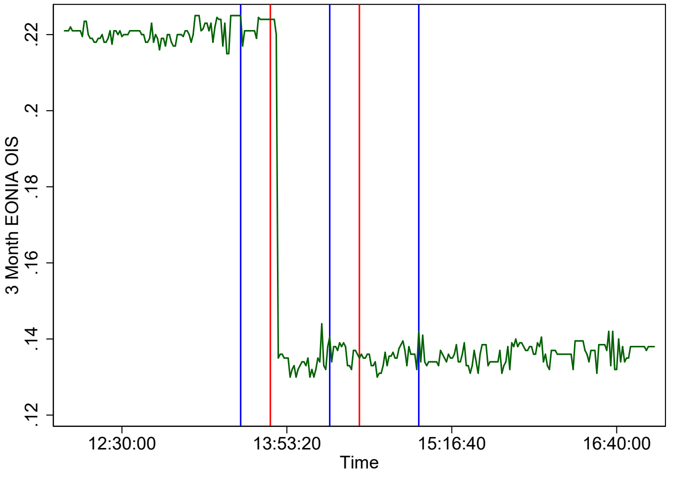

Figure 4.1: *Overnight Indexed Swap rates on 05.07.2012*

Note: This figure shows the time series of the 3 month EONIA Overnight Indexed Swap rate for July 5, 2012. The blue lines represent the borders of our measurement windows, the red lines indicate the policy events, i.e., the ECB's press release at 13:45 CET and the start of the press conference at 14:30 CET. The first pre-window runs from 13:30-13:44 CET and then the first post-window is active between 13:45-14:14 CET. Then a second pre-window runs from 14:15-14:29 CET and the second post-window is active between 14:30-14:59 CET.

We use the 3 month OIS rate obtained from Datascope. To convert the instrument series obtained in this way to monthly frequency, we follow Gertler and Karadi (2015). Because the announcement days are at different times during each month, we weigh each observation according to when in a month it occurred. Let a_d be the cumulative shock series at day d in the month, which evolves in the following way

$$a_d = \begin{cases} a_{d-1} + \Delta f_d & \text{if announcement at day } d \\ a_{d-1} & \text{otherwise} \end{cases}$$

where Δf_d is the change in the OIS rates calculated as described above. We then weight the series according to

$$F_t = \frac{1}{D_m} \sum_{d \in m} a_d \tag{4.1}$$

where D_m is the number of days in month m. Finally, the instrument for each month t is

$$Z_t = F_t - F_{t-1} \tag{4.2}$$

4.2.2 The effect of monetary policy shocks on output

We follow Jordà (2005) and Stock and Watson (2018) and estimate the response of output to monetary policy shocks using the local projections instrumental variable (LPIV) method, employing the instrument discussed in the previous section. Impulse responses, for each country n, are constructed from the sequence $\{\beta_n^h\}_{h=0}^H$ from the following equations

$$y_{n,t+h} - y_{n,t-1} = \alpha_n^h + \beta_n^h \hat{i}_t + \sum_{j=1}^{p} \Gamma_{n,j}^h X_{t-j} + u_{n,t+h}, \quad h = 0, \ldots, H \tag{4.3}$$

where y_n is log of output in country n, X is a set of control variables and $\hat{i}$ are the fitted values from the first-stage regression[10]

$$i_t = c + \rho Z_t + \sum_{j=1}^{p} D_j^h X_{t-j} + e_t \tag{4.4}$$

As a benchmark we set the number of lags to $p = 3$ and the horizon of the impulse responses to $H = 36$.[11] In all specifications we include the interest rate i, the instrument Z, aggregate euro area output and the price level in our set of control variables, X.[12,13] Notice that Equation (4.4) resembles a standard Taylor rule for the ECB: the current interest rate depends on lags of euro area output and inflation, plus

[10]For a detailed description of the data series used, we refer the reader to Appendix 4.H.

[11]We also estimate specifications in which the number of lags is allowed to vary across countries using the *Akaike Information Criteria* (AIC). Doing so leaves the results unaltered and therefore, for simplicity, we choose the same number of lags for all countries.

[12]As pointed out by Ramey (2016), the construction of the instrument as in Gertler and Karadi (2015) introduces auto-correlation into the instrument, invalidating our identifying assumptions. To alleviate this problem, we include the instrument in addition to the other control variables in X.

[13]Removing the set of lagged control variables in (4.4), specially the interest rate, leads to very low F-statistics. Since the interest rate is persistent, contemporaneous shocks account for only a small part of its variance. Furthermore, of the contemporary shocks, the monetary policy shock is only a fraction. Therefore, the explanatory power of the instrument alone on the interest rate can be expected to be fairly low (Stock and Watson, 2018). The first state F-statistic in our benchmark specification is 17.42.

lags on the interest rate itself.[14]

Our dependent variable, monthly GDP, is measured as the logarithm of real GDP. Given that GDP is only available at quarterly frequency we follow Chow and Lin (1971) to interpolate real GDP into a monthly frequency.[15] We use the Euro Overnight Index Average (EONIA) as the monetary policy rate and the logarithm of the deseasonalized Harmonized Index of Consumer Prices (HICP) as the measure of the aggregate price level. We use data from January 2000 to December 2012, capturing the initial stages of the adoption of the euro and ending during the year when the interest rate hit the zero lower bound.

Figure 4.2 presents impulse responses of real GDP for each country in our sample to an expansionary shock of one standard deviation in our instrument, following Jarociński and Karadi (2020). The IRFs are represented by the blue lines and, following Stock and Watson (2018), we construct 1 and 2 standard deviation confidence bands that surround the point estimates, using Newey-West estimators.[16]

The estimated impulse response functions reveal that expansionary monetary policy shocks cause output to increase in most countries. Output increases significantly after a little less than a year, with the maximum impact most often occurring later. The response of aggregate euro area output, for example, reaches its peak of 0.23 percentage points after 27 months.

[14] As a robustness check, we conduct the same exercise, including country-specific lags in Equation (4.4), leading to country specific first-stage regressions and country specific $\hat{\imath}$ s. The results are reported in Figure 4.11 in the Appendix.

[15] For each euro area country, as well as the aggregate euro area, we use monthly data for industrial production, retail trade and unemployment to construct monthly series for real GDP.

[16] The local projection impulse responses for prices are presented in Figure 4.12 in the Appendix.

There is considerable heterogeneity in the magnitude of the responses, both in peak and cumulative effect. Moreover, the initial impacts of monetary policy shocks seem to be on average small and often not statistically different from zero.

Using the results from Georgiadis (2015) as a proxy for the VAR counterpart of our analysis, we find that the peak values are strongly correlated for the subset of countries that overlap with his sample, with a correlation coefficient of approximately 0.84.[17] Given that the relative positions of countries is important for the analysis in the upcoming section, we consider it to be reassuring that our estimates are in line with the findings in Georgiadis (2015).

We proceed now by relating the strength of the output responses to the share of liquidity constrained households in each country.

[17] Georgiadis (2015) estimates responses for Austria, Belgium, Finland, France, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Slovakia, Slovenia and Spain.

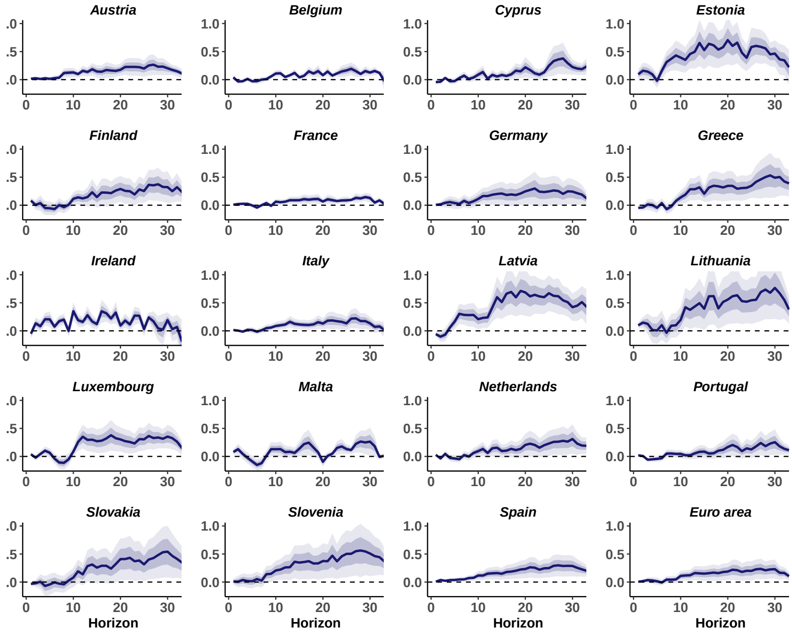

Figure 4.2: *Impulse responses for output in euro area countries – LPIV*

Note: This figure shows impulse responses of real GDP to an expansionary monetary policy shock of one standard deviation. For each euro area country, the response is estimated using LPIV (Equation 4.3). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark and light blue shaded areas represent 1 and 2 standard deviation confidence bands, respectively, constructed using Newey-West estimators.

4.3 Measuring Financial Constraints

Bilbiie (2020) describes a TANK economy in which household heterogeneity is collapsed to being either financially constrained or not. Taking this idea to the data, our aim is to construct variables that measure the degree of financial constraints in a given country. To do so, we rely on the Eurosystem Household Finance and Consumption Survey (HFCS) and the European Union Statistics on Income and Living Conditions (EU–SILC). In the next subsections, we describe these datasets and the construction of our measures for financial constraints used in the subsequent analysis.

4.3.1 The Household Finance and Consumption Survey

The HFCS is conducted by the Household Finance and Consumption Network (HFCN), tasked by the Governing Council of the ECB. The survey is modeled after the US Survey of Consumer Finances and is harmonized across the euro area, set up to collect micro data on household finances (Honkkila and Kavonius, 2013). It contains data from interviews with over 84,000 households. Three waves have been conducted, with data releases in 2013, 2016 and 2020. In our main analysis, we rely on data from the second wave.

In approximating the share of households who have high MPC, we follow Kaplan et al. (2014).[18] A household is categorized as living HtM if its liquid wealth is smaller than a certain share of monthly income. In their set of countries, the share of HtM households is between 20 to 35 percent Kaplan et al. (2014).[19]

[18] Kaplan et al. (2014) find that households categorized as HtM according to their measure have an estimated MPC of more than twice that of non-HtM households.

[19] U.S., Canada, Australia, U.K., Germany, France, Italy, Spain.

Let m_i denote liquid assets, a_i denote illiquid assets, y_i denote income and $\underline{m}_i$ be a credit limit for household i.[20] We categorize a household as HtM if:

$$0 \leqslant m_i \leqslant \frac{y_i}{2} \tag{4.5}$$

or if:

$$m_i \leqslant 0, \quad \text{and} \quad m_i \leqslant \frac{y_i}{2} - \underline{m}_i \tag{4.6}$$

The credit limit $\underline{m}_i$ is set to be the household's monthly income, capturing the possibility of spending using a credit card and repaying the debt once a month. For our sample, the fraction of households who are categorized according to Equation (4.6) is small.

We further divide households into wealthy and poor HtM (Kaplan et al., 2014). A household is categorized as wealthy HtM if, in addition to fulfilling one of the conditions in Equations (4.5) and (4.6), it has positive illiquid wealth:

$$a_i > 0 \tag{4.7}$$

If a household satisfies either one of Equations (4.5) or (4.6), but not the condition in Equation (4.7), we label that household as poor HtM.[21]

Figure 4.3a plots the total fraction of HtM households as well as the split between wealthy and poor. The cross-country variation is striking, with the fraction of HtM households ranging from just

[20] See Appendix 4.F for details about the classification of assets as liquid and illiquid.

[21] For a discussion about the theory behind this classification scheme, we refer the reader to Kaplan et al. (2014).

above 10 percent in Malta to almost 65 percent in Latvia. In most countries (exceptions being Austria, France, Germany and Ireland) the fraction of wealthy HtM households exceeds the fraction of poor HtM households, which is in line with the findings in Kaplan et al. (2014).[22]

A concern with the measure described above is that households are interviewed at different points during the month or the year. If there are systematical differences across countries in when households are interviewed, this could lead to biased estimates. To combat this, we construct a second proxy for a household's MPC which relies on their past year's income and expenses.

In the HFCS questionnaire, households are asked if, over the last 12 months, their expenses (i) exceeded income, (ii) were about the same as income or (iii) were less than income. A household in categories (i) or (ii) is likely more sensitive to unexpected shocks than one in category (iii), and is, therefore, likely to have a higher MPC. We compute the fraction of households whose expenses were about the same as or exceeded income (categories (i) and (ii)) and label households that fulfill this criterion as being "Potentially Financially Vulnerable type 1" (PFV1).

The fractions are presented in Figure 4.3b. Again, there is heterogeneity across countries and the average, indicated by the vertical line, is above 60%. For all countries, the share of PFV1 households is higher than the HtM share. We consider this statistic an upper bound for the fraction of households who have high MPC, as it disregards

[22]Most households that are classified as W-HtM own the residence in which they live. The data show that in most countries the majority of households that are classified as W-HtM do not have a mortgage; the fraction varies between 0.12 and 0.67 with mean (median) of 0.34 (0.34). This fraction appears to be negatively correlated with the fractions of W-HtM across countries. See Appendix 4.G.1 for more details

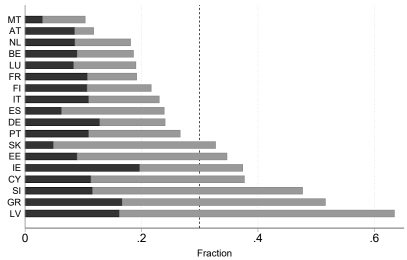

(a)

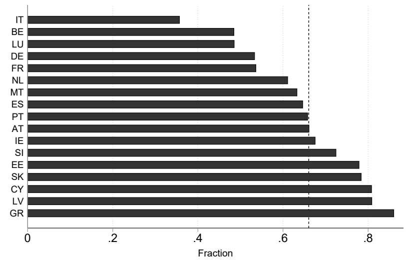

(b)

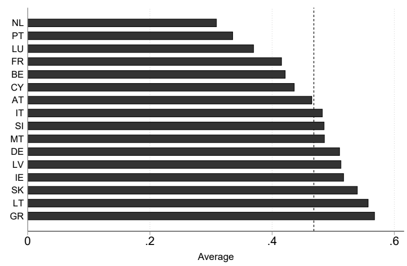

(c)

Figure 4.3: *HFCS proxies for Hand-to-Mouth status*

Note: *Panel (a)*: This figure shows the fraction of households that are classified as HtM in each country. The total fraction, given by the total length of each bar, is divided up into two parts: poor (black) and wealthy (gray). The vertical line indicates the unweighted average of total HtM in our sample of countries. We do not have data for Lithuania. *Panel (b)*: The figure shows the fraction of households that have had expenses over the last 12 months that were "about the same as" or exceeded their income over the same period. The total fraction is given by the total length of each bar. The vertical line indicates the average of the fractions in our sample of countries. Data is missing for Finland and Lithuania, hence they are not included in the figure. *Panel (c)*: The figure shows the average of fractions of lottery winnings, in each respective country, that the households would spend over the next 12 months. See text for a more detailed description. Data does not exist for Estonia, Finland and Spain. For panels (a) and (b), we use data from the second wave of the Eurosystem Household Finance and Consumption Survey (HFCS). For panel (c), we use data from the third wave of the HFCS.

the possibility that they might have substantial amounts of liquid assets. The correlation between PFV1 and HtM is 0.68, which we see as encouraging.

The most recent wave of the HFCS introduces a new question which attempts to capture MPC in a more direct way. Households are asked what percentage of a hypothetical lottery win they would spend over the next 12 months.[23] Within each country, we compute the average of these reported MPC across all households. Figure 4.3c presents the resulting averages and we can see that there is considerable variation across the countries. The correlation between this measure and our HtM measure is 0.49.

4.3.2 The European Union Statistics on Income and Living Conditions survey

The sample period for the two measures derived above coincides with the end of the European Sovereign Debt Crisis. To ensure that this is not driving our results, we construct two additional variables from the European Union Statistics on Income and Living Conditions (EU-SILC) questionnaire. The EU-SILC is a yearly survey with the objective to measure income, poverty, social exclusion and living conditions in the European Union and is executed by the national statistical authorities. At its introduction in 2003 it covered seven countries, and since 2005 has covered all the countries in our sample, with a sample size of close to 90,000 households.[24] Because of its early inception,

[23]The question reads: "Imagine you unexpectedly receive money from a lottery, equal to the amount of income your household receives in a month. What percent would you spend over the next 12 months on goods and services, as opposed to any amount you would save for later or use to repay loans?"

[24]The sample size for the whole survey is about 130,000 households. The figure in the text refers to the 19 countries in our sample. The country specific averages

the survey allows us to construct proxies for the share of households with high MPC with data from before the Great Recession.

First, we use a question on whether a household, out of its own resources, would be able to cover a hypothetical, unexpected, required financial expense, equal to the national monthly at-risk-of-poverty threshold.[25] Households who expect not to be able to do so are likely to have high MPC out of transitory income shocks. We take the share of households unable to face an unexpected expense as a percentage of all households in 2005 and label it "Potentially Financially Vulnerable type 2" (PFV2). Figure 4.4a displays the variable across countries. Although it is calculated using a different survey and a different sampling period, the correlation coefficient between PFV2 and HtM is 0.67.

We construct one more variable using the EU-SILC survey from 2005. In the survey, households are asked if they were unable to pay utility bills during the last year on time (have been in arrears) due to financial difficulties.[26] We assume that households to whom this applies will consume a large share of an unexpected income shock and therefore classify these households as having high MPC and all others as having low MPC. The share of the former in the population is "Potentially Financially Vulnerable type 3" (PFV3). The cross-sectional distribution of PFV3 is shown in Figure 4.4b. For all countries, this measure is the lowest. Intuitively, all other indicators measure the potential of not being able to "make ends meet" for a household, while PFV3 is the share of households who are already behind on making payments. Therefore, it can be viewed as the strictest proxy among

produced from the SILC-EU survey are obtained from Eurostat.

[25]The at-risk-of-poverty threshold is defined as 60% of the national median equivalized disposable income after social transfers.

[26]Utility bills include heating, electricity, gas, water, etc.

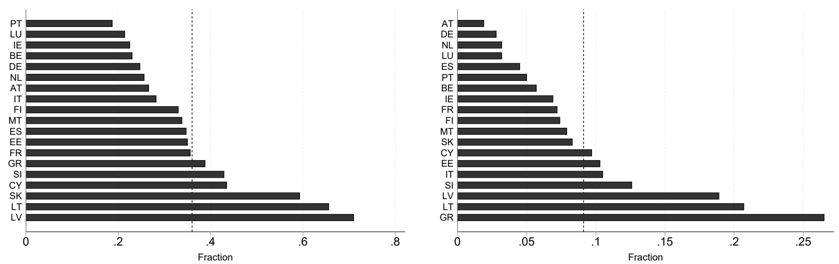

(a) ***(b)***

Figure 4.4: *PFV2 and PFV3*

Note: *Panel (a)*: The figure shows the fraction of households that believe that they are unable to face unexpected expenses with the use of own resources (PFV2). The fraction is given by the length of each bar. The vertical line indicates the average of the fractions in our sample of countries. The data is from European Union Statistics on Income and Living Conditions (EU-SILC), obtained from Eurostat (2019e). *Panel (b)*: The figure shows the fraction of households that, over the last year, were in arrears on their utility bill (PFV3). The fraction is given by the length of each bar. The vertical line indicates the average of the fractions in our sample of countries. Data is from European Union Statistics on Income and Living Conditions (EU-SILC), obtained from Eurostat (2019b).

the ones presented in this section and we view it as the lower bound of households with high MPC. The correlation between PFV3 and the HtM measure is 0.80.

4.4 Liquidity Constrained Households and Monetary Policy Effectiveness

4.4.1 Results

The results in Section 4.2.2 indicate that the countries in our sample do not respond homogeneously to monetary policy shocks. We proceed to link this finding with country-specific aggregates, which relate to asset- and income positions of households. Our primary focus is on the share of households living HtM, but we also report results for the three alternative measures introduced in Sections 4.3.1 and 4.3.2: PFV1, Lottery MPC, PFV2 and PFV3.

Scatterplots between different measures of monetary policy effectiveness and the shares of households living HtM are presented in Figure 4.5. Both panels show the share of HtM households on the horizontal axis and the vertical axes display different measures of the effectiveness of monetary policy. Figure 4.5a shows the peak of the output impulse response, which exhibits a significant positive correlation with the HtM share across countries.[27] Figure 4.5b instead uses the cumulative impulse response, with very similar results. Both suggest that an accommodating monetary policy shock has bigger effects on output in countries with a higher share of HtM households.

[27]Because we calculate both the peak responses and the HtM shares, there is uncertainty associated with our point estimates. In order to not clutter the figure reported here, we relegate the scatterplot including confidence intervals to Figure 4.10a in the Appendix.

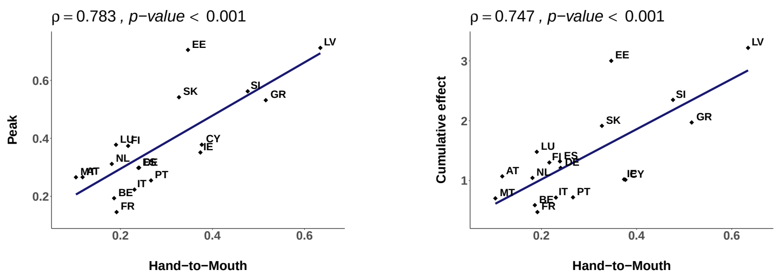

(a) *(b)*

Figure 4.5: *Monetary policy effectiveness and Hand-to-Mouth shares*

Note: This figure plots the effectiveness of monetary policy, as measured by the peak effect and cumulative effect of the real GDP impulse responses, calculated using the benchmark LPIV estimation, against the share of households classified as living HtM in each euro area country (except Lithuania, not included in the HFCS). The HtM shares are calculated using data from the Eurosystem Household Finance and Consumption Survey. The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of Peak/Cumulative values on HtM shares. In the upper left corner of each panel, we report the correlation coefficient ρ and the p-value. *Panel (a)*: Peak effects and share of Hand-to-Mouth. *Panel (b)*: Cumulative effects and share of HtM, normalized by aggregate euro area cumulative effect.

We interpret these results in light of a standard TANK model as in Bilbiie (2020). Here, a certain fraction of households consume their income every period, by construction, while the remainder can save and borrow. This simple setup captures an important element of monetary policy transmission with heterogeneous agents: a partial and a general equilibrium effect. The former describes output effects which occur due to the Euler equation of the unconstrained households. A shock which lowers the real interest rate makes these households demand more output in the current period. The general equilibrium effect includes the changes in output caused by changes in wages and profits. In the model, whether the share of constrained households amplifies (dampens) the aggregate output response depends on whether income, i.e., the sum of wage and profit income, of constrained households moves more (less) than one-for-one with aggregate income.

The results in Figure 4.5 show that a higher share of liquidity constrained households amplifies an economy's response to monetary policy surprises. As explained above, this is in line with Bilbiie (2020) if the income elasticity of constrained households with regard to aggregate income is larger than one. Richer models such as those in Auclert (2019) or Hagedorn et al. (2019) feature more channels through which different households can be differently affected by aggregate shocks; still, they imply that if the income of the highest MPC agents comoves more with aggregate income than that of the low MPC agents, this mechanism amplifies the economy's response to shocks relative to RANK models.

Using the panel dimension of the HFCS, we find suggestive evidence that the elasticity of HtM households' three year income growth with respect to aggregate (three year) income growth is significantly higher than that of non-HtM households. For details, see Appendix 4.A. These findings are in line with Patterson (2019), who, in a US context, provides evidence for MPC being larger for individuals who are more affected by business cycles.

We now turn to our alternative measures of MPC, namely PFV1-PFV3 and the self-reported propensity to consume out of lottery winnings. Our focus is on peak responses, but as before, results are similar using cumulative responses as the measure for monetary policy effectiveness.

The four scatterplots are presented in Figure 4.6. Correlations between the peak responses and each of the four statistics are strong. We view this as encouraging for two reasons. First, the results lend credence to the measure proposed by Kaplan et al. (2014). The correlations are very similar, although the alternative proxies use different approaches and, in two cases, different surveys. Second, the proxies for MPC in panels (c) and (d) were calculated using data from 2005, giving us confidence that our results are not driven by the Financial Crisis or the European sovereign debt crisis. On the contrary, the correlations we find are a persistent feature of European monetary policy transmission.

Next, we investigate the importance of the distinction between liquid and illiquid asset holdings in more detail. Kaplan and Violante (2014) and Kaplan et al. (2014) argue that it is important to disaggregate these two types of assets by partitioning households into Wealthy HtM households (liquidity constrained but owning positive illiquid wealth) and Poor HtM households (zero or negative illiquid wealth). They estimate the MPC of P-HtM (W-HtM) households to be twice (thrice) as large as the MPC of unconstrained households. However, a sole focus on differences between P-HtM and W-HtM

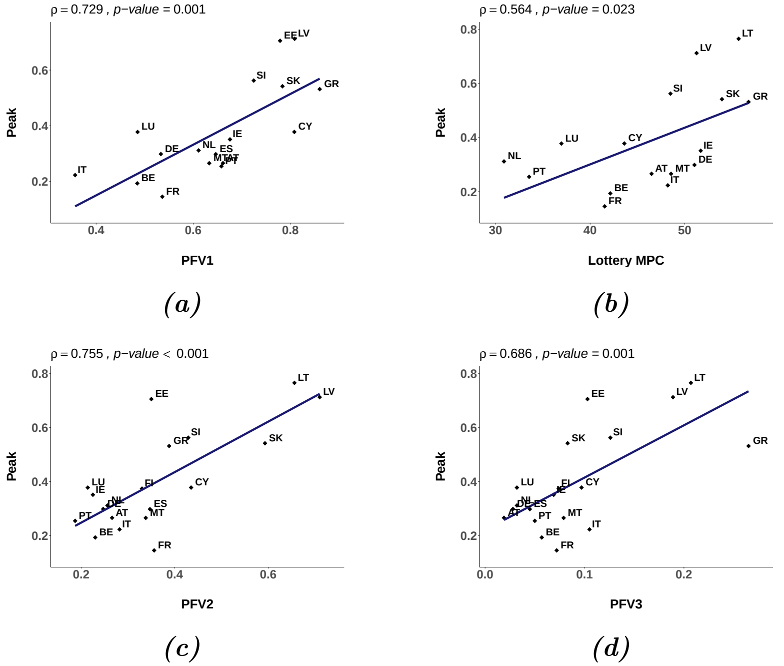

Figure 4.6: *Impact of monetary policy and alternative liquidity constraint measures*

Note: This figure plots the effectiveness of monetary policy, as measured by the peak effect, calculated using the benchmark LPIV estimation, against different statistics in each euro area country. The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of peak values on the respective statistics. In the upper left corner of each panel, we report the correlation coefficient ρ and the p-value. *Panel (a)*: Peak effects and PFV1. *Panel (b)*: Peak effects and Lottery MPC, calculated using data from the Eurosystem Household Finance and Consumption Survey (HFCS). *Panel (c)*: Peak effects and PFV2. *Panel (d)*: Peak effects and PFV3. PFV1-PFV3 and Lottery MPC as defined in Section 4.3. PFV1 is calculated using data from the HFCS and to calculate PFV2 and PFV3 we use data from European Union Statistics on Income and Living Conditions (EU-SILC).

households in MPC overlooks that their incomes might adjust differently and a potential revaluation of illiquid asset portfolios of W-HtM households, following a shock to the interest rate Auclert (2019). Our data does not allow us to investigate how income and asset values change following the shocks.

Figure 4.7 shows the relationship between the peak responses across countries and their shares of wealthy and poor HtM households, in panels (a) and (b), respectively. While the W-HtM share is strongly correlated with the peak values of the IRFs, this is not the case for the share of P-HtM households. This suggests that disregarding households' liquidity positions, in theoretical models and empirical work, can lead to erroneous conclusions about the effects of monetary policy, as argued by Kaplan et al. (2014). We view this as an interesting question for future research.

As a complementary test, we investigate the relationship between peak responses and the fraction of asset poor households. Kaplan et al. (2014) argue that total net wealth, which is the standard metric for high MPC behavior in heterogeneous-agent macroeconomic models, is a poor predictor of MPC. As in Kaplan et al. (2014), a household is labeled as asset poor if the sum of its net wealth is zero or negative. Panel (c) in Figure 4.7 gives no evidence for a relationship between output responses and the share of asset poor. The statistic is outperformed by all of our other measures in predicting by how much output is affected through monetary policy shocks.

4.4.2 Robustness

First, we test whether our results are affected by restricting the sample to the countries who adopted the Euro by the year 2002, when the currency was introduced. For this set of countries, the ECB was

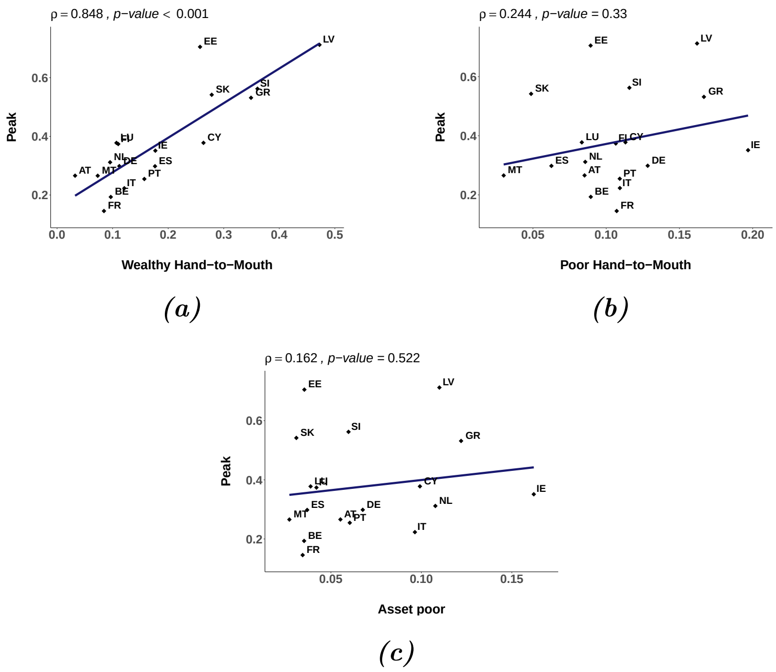

Figure 4.7: *Monetary policy effectiveness and Wealthy and Poor Hand-to-Mouth shares*

Note: This figure plots the effectiveness of monetary policy, as measured by the peak effect, calculated using the benchmark LPIV estimation, against the share of households classified as living as Wealthy HtM, Poor HtM and asset poor Kaplan et al. (2014), respectively, in each euro area country (except Lithuania). The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of peak values on Wealthy HtM shares, Poor HtM shares and the share of asset poor, respectively. In the upper left corner of each panel, we report the correlation coefficient ρ and the p-value. *Panel (a)*: Peak effects and share of Wealthy Hand-to-Mouth. *Panel (b)*: Peak effects and share of Poor HtM. *Panel (c)*: Peak effects and asset poor. Wealthy HtM shares, Poor HtM shares and shares of assets poor are calculated using data from the Eurosystem Household Finance and Consumption Survey.

the relevant monetary policy institution throughout our sample period. The first column of Table 4.1, for reference, reports the results outlined in the previous section (reported in Figure 4.5a). The second column reports the same statistics for the sample of initial euro area members. Although the correlation between the share of HtM households and the peaks of the IRFs is attenuated slightly, it is still 0.7 and statistically significant. We see this as encouraging, as the conclusions drawn in the previous section are not driven by countries which joined the currency union after 2002.

Table 4.1: *Robustness tests*

	(1) Baseline	(2) Init Members	(3) 1st wave	(4) Consumption	(5) Cons - IM	(6) GVAR	(7) JK
ρ	0.78	0.70	0.69	0.75	0.88	0.78	0.58
t-statistic	5.04	3.07	2.87	4.43	5.84	4.91	2.86
p-value	0.00	0.01	0.02	0.00	0.01	0.00	0.01

The table shows the correlation coefficient ρ between a measure of the share of HtM households and the peak output response to a monetary policy shock across countries. The second and third row display the associated t-statistic and p-value, respectively. The *first column* shows the results for our baseline specification. The *second column* shows the results when we restrict the sample to euro area countries which were members in 2002. The *third column* restricts the sample to early euro area members using the first wave of the HFCS to compute HtM shares. The *fourth column* uses the peak response from a *consumption* IRF, as opposed to GDP. The *fifth* column uses peak responses from a consumption IRF and restricts the sample to early euro area countries. The *sixth* column obtains the peak responses from a GVAR outlined in Appendix 4.C. The *seventh* and last column uses peak responses produced using the Monetary Policy shocks series in Jarociński and Karadi (2020).

Along similar lines, we can test whether using the first wave of the HFCS, conducted in 2010, affects our conclusions. Column 3 in Table 4.1 reports the correlation between HtM shares computed from the HFCS' first wave and peak responses of GDP after a monetary policy surprise. Importantly, during the first wave, the HFCS was not conducted in all countries in our sample, which leads us to restrict the analysis to the initial members of the euro area. The relevant comparison, hence, is the second column in table 4.1. While the t-statistic becomes slightly smaller, the point estimates are almost equivalent across different survey waves.[28]

Because consumption, as opposed to GDP, is the relevant metric for household welfare, columns 4 and 5 in Table 4.1 repeat the exercise from section 4.4.1, substituting GDP with a quarterly measure of household consumption. As before, we interpolate it to monthly frequency.[29] The results for the full sample (column 3) are very similar to those estimated using the monthly GDP series. The correlation coefficient falls very slightly from 0.78 to 0.75. Restricting the sample to the initial members of the euro area (column 5), the correlation coefficient increases considerably to 0.88. The scatterplots associated with these estimations are reported in Figure 4.8.[30] These results indicate that our initial findings are not driven by heterogeneous investment demand or fiscal responses across countries.[31]

As another robustness check, we construct a GVAR for the euro

[28]We report the fractions of HtM households across countries according to all three survey waves of the HFCS in Figure 4.13 in the Appendix. The fractions are remarkably stable across time.

[29]For each euro area country we again use monthly data for industrial production, retail trade and unemployment to construct monthly series household consumption.

[30]See Figure 4.10b for the associated scatterplot including confidence bands.

[31]This conclusion assumes $\mathsf{GDP} = \mathsf{C} + \mathsf{I} + \mathsf{G}$.

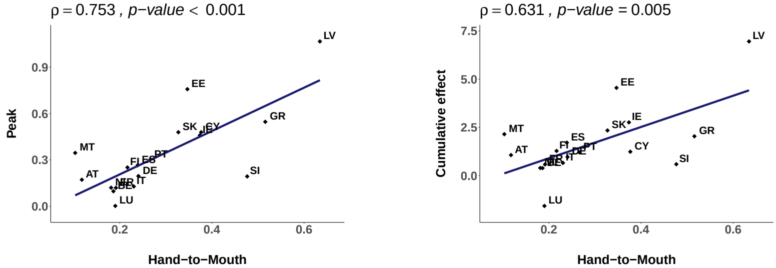

(a) ***(b)***

Figure 4.8: *Monetary policy effectiveness and Hand-to-Mouth shares – consumption responses*

Note: This figure plots the effectiveness of monetary policy, as measured by the peak effect and cumulative effect of the total household consumption impulse responses, calculated using the benchmark LPIV estimation, against the share of households classified as living HtM in each euro area country (except Lithuania, not included in the HFCS). The HtM shares are calculated using data from the Eurosystem Household Finance and Consumption Survey. The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of Peak/Cumulative values on HtM shares. In the upper left corner of each panel, we report the correlation coefficient ρ and the p-value. *Panel (a)*: Peak effects and share of Hand-to-Mouth. *Panel (b)*: Cumulative effects and share of HtM.

area based on the work by Georgiadis (2015) and Burriel and Galesi (2018), and repeat our analysis in this framework. See Appendix 4.C for details on the model setup. We utilize the same data as for the LPIV estimation and to our knowledge, we are the first to combine the instrumental VAR techniques laid out in Stock and Watson (2018) with the GVAR setting. The correlation coefficient between the peak responses estimated using this approach and the HtM shares across countries is reported in column 6 of Table 4.1. It is identical to the same statistic obtained from the LPIV estimation, and highly significant.

Lastly, we show that our results are robust to using the shock-series produced in Jarociński and Karadi (2020), who distinguish between monetary policy shocks and information shocks. We use the former series and repeat the analysis above, constructing new impulse response functions and obtaining new peak values.[32] Column 7 in Table 4.1 shows that the correlation statistic between the share of HtM households and the peak responses is lower than it is with our shock series, but still highly significant.

Next, we show that our results are robust to changing the horizon at which the effects of monetary policy are measured. In the previous section, we mainly rely on peak values. Here we instead focus on the point estimates at different horizons $h = \{0, 1, \ldots H\}$ and first extract the point estimate for each country n, β_n^h, to then correlate each of these with the HtM values.

The horizontal axis in Figure 4.9a shows the horizon (h) and the vertical axis shows the correlation between the country specific HtM measures and IRF point estimates. During the majority of the first

[32]The resulting impulse responses for output and prices are reported in Appendix 4.B.

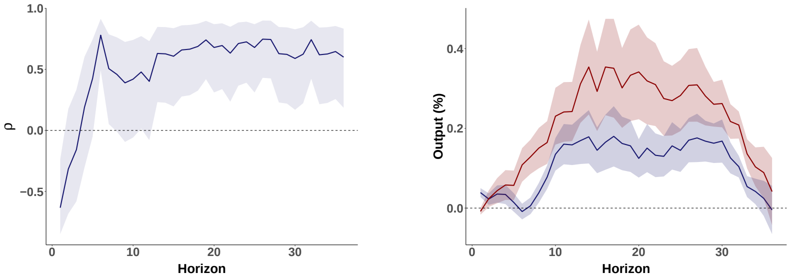

(a) ***(b)***

Figure 4.9: *Monetary policy effectiveness and Wealthy and Poor Hand-to-Mouth shares*

Note: *Panel (a)*: correlation between HtM and responses at different horizons, retrieved from the impulse responses using the benchmark specification, for each horizon $h = \{0, 1, \ldots, H\}$. The shaded area is the 95% confidence band around the point estimates. *Panel (b)*: IRFs for two groups of countries. The blue line represents the IRF for the group consisting of countries with HtM shares below the median and the red line represents the IRF for the group consisting of countries with HtM shares above the median. See text for details. The shaded areas give 68% confidence bands around the point estimates. Calculations of HtM shares are based on data from the Eurosystem Household Finance and Consumption Survey.

year, the correlation is not significant. This is unsurprising, as monetary policy affects output with a lag. After a year, however, the correlation is statistically and economically significant until it dies out towards the end of our estimation horizon. The latter, again, is unsurprising, as Figure 4.2 indicates that the effect of a common monetary policy shock peters out after three years in most countries.

Second, we divide the countries into two groups based on HtM shares; countries with HtM shares below the median are placed in the first group and countries with HtM shares above the median are placed in the second group. We then re-estimate Equation (4.3) for each of the two groups.[33]

Figure 4.9b graphs the results for the two group specific IRFs. We again find that output reacts more to monetary policy shocks in countries with higher HtM shares. We view these results as strengthening our previous conclusion that the share of HtM households is a relevant statistic for the effectiveness of monetary policy across countries.[34]

Next, we investigate whether other country-specific characteristics can account for the heterogeneity in impulse responses that we observe. We focus on a set of variables that could be correlated with both HtM shares and the effectiveness of monetary policy. Our strategy is the following: First, we gather data on variables suggested in the literature as relevant for the effectiveness of monetary policy in a cross country perspective. Subsequently, for each variable, we investigate (i) whether it is correlated with the HtM shares, (ii) whether it is correlated with the peak effects we find in section 4.2 and (iii)

[33] After dividing countries into the two groups, we then index the GDP series of each country and use the average index values within each group as a measure for GDP.

[34] We can perform the same analysis using a Panel IV setup, including country fixed effects. This approach is discussed in Appendix 4.D, where we show that the inclusion of such fixed effects does not change our conclusions.

whether after controlling for the variable, the HtM shares still explain a part of the output responses we observe.[35]

Cloyne et al. (2020) find that households who own a mortgaged property adjust consumption spending more than both renters and homeowners without mortgages, in response to unexpected interest rate changes. The authors find that consumption among homeowners without mortgages is insensitive to changes in monetary policy. It is furthermore possible that monetary policy can affect house prices and output via the collateral channel (see Cloyne et al., 2019). Corsetti et al. (2018) find that the strength of the housing channel is related to home ownership rates. These results lead us to the first three variables that are introduced in this section. The first variable is labelled Own and represents the fraction of households in each country that own their main residence. We allow for an outstanding mortgage to be tied to the residence. A closely related variable is Mort, which represents the fraction of households in each country that have a mortgage. Additionally, in each country there are households that own their main residence but do not have a mortgage attached to it. We label the variable for the fraction of these households in each country as HO.

It is possible that the effectiveness of monetary policy depends on how highly indebted households are (see, e.g., Flodén et al., 2020) and on how common it is that mortgages have an adjustable interest rate (see, e.g., Calza et al., 2013; Flodén et al., 2020). We calculate the fraction of households that have *at least one* mortgage with an adjustable interest rate and label the variable Flex. To test if HtM shares and effectiveness of monetary policy are related to how highly

[35] Most of these control variables are constructed using data from the HFCS, since many of them are related to housing and how it is financed.

indebted households are, we calculate average loan-to-value ratios and average loan-to-income ratios among households with mortgages in each country and label them LTV and LTI, respectively. Households with LTV ratios above 1.5 were removed in the calculations of LTV and observations with LTI ratios above 10 were removed from the calculations of LTI, to limit the influence of outliers.

In Section 4.4.1 we argue that it is mainly the fraction of wealthy HtM households that explains why the total fraction of HtM households is correlated with peak values. It is possible that this result is driven by the share of households with positive amounts of illiquid wealth, not necessarily by the share of HtM. We can rule this out by showing that the correlation between wealth HtM shares and the peak values is not driven by the shares of wealthy households across countries. To this end, we calculate the share of wealthy households in each country and label the variable $\mathsf{Wealthy}$.[36]

Wong (2019) finds that, in the U.S., especially younger households refinance loans following changes in the interest rate and drive most of the aggregate response in consumption. Examining the HFCS data, we observe that older households, on average, are less likely to be HtM. Moreover, the probability of being wealthy HtM increases between age 20 and the late 30s, and decreases after this threshold. We find it important to test if including the average age in each country as a control variable changes our results. The average age of household heads in our data is labelled Age.

The growing literature using GVAR models (e.g., Burriel and Galesi, 2018; Georgiadis, 2015) emphasizes the importance of considering spillover effects of monetary policy, with the sizes of these spillovers partly related to trade flows. We include a measure of trade

[36] A household whose net illiquid assets are positive is labelled as wealthy.

openness due to its importance in the Dynamic IS equation in the small open economy literature (Galí and Monacelli, 2008). We calculate trade openness as the sum of imports and exports as a share of GDP in each country, to test if what we find is related to trade. We use the World Bank national accounts data to calculate this statistic and label it Trade.

The next variable is labelled ROL and is related to how regulated labor markets are. Georgiadis (2015), using data from a subset of the countries that we consider, estimates that output in countries with more regulation respond less to monetary policy shocks. We construct this variable by calculating the average of the "Employment laws index" and the "Collective relations laws index" from Botero et al. (2004). Georgiadis (2015) also finds that the share of GDP accounted for by services is closely connected to the effectiveness of monetary policy, showing that countries that have the lowest shares, compared to countries with the highest shares, exhibit responses of output which are half as large. We have mentioned that our estimates for the effectiveness of monetary policy are similar to the estimates in his paper. Hence it is likely that our measures are also correlated with service shares and it becomes important to see if there is variation left in HtM shares, even after having controlled for service shares, that is correlated with the effectiveness of monetary policy. We label the variable $\mathsf{Service}$. To calculate it we average over the shares reported in the World Bank (2019) WDI database between years 2000-2012 for each country.

Our sample period coincides with large house-price fluctuations in some European countries. In order to show that the sizes of our HtM shares are uncorrelated with these changes, we control for a measure of house price growth across European countries. We utilize

Eurostat's house price index, which starts in 2005. House price growth is calculated as the average quarterly year-on-year change in the index between the first quarter for which data are available and the last quarter of 2012.[37] We label the variable *HP Growth*.

Economic development is potentially correlated with how countries respond to shocks and with the share of HtM households. For this reason, we control for GDP per capita of 2008. We label the variable as *GDPpc*.

All results are summarized in Table 4.2. The first column in the table presents raw correlations between the peak effects and the different variables that vary across the rows in the table. In the second column we see the correlations between the HtM shares and the variables that vary across the rows. Most often the absolute values of the correlation coefficients are relatively close to zero. One exception is HO for which the correlation is positive and of significant magnitude for both peak effects and HtM shares.[38] Another is $\mathsf{Services}$ which is negatively correlated with peak effects (confirming the result from Georgiadis (2015)) and also negatively correlated with HtM shares.

That fact that peak effects and HtM shares are correlated with some of these variables was expected. The important question is whether these other variables are likely to be the reason we find such a strong correlation between peak responses and HtM shares. To get a sense of whether this could be the case, we calculate semipartial correlations between the estimated peak effects and HtM shares. These

[37]For most countries, the first data point is available in 2005. Data for Italy and Austria is only available since 2010. The index is not available for Greece

[38]The negative correlation between Mort and HtM might seem surprising since it appears plausible that Wealthy HtM households often have mortgages. In appendix 4.G.1 we show this to not be the case. Another potentially surprising finding, given results in Cloyne et al. (2020), is the positive correlation between Peak and HO. We investigate and discuss it further in appendix 4.G.2

Table 4.2: *Correlations and semipartial correlations*

X	$\rho(\mathsf{Peak}, \mathsf{X})$		$\rho(\mathsf{HtM}, \mathsf{X})$		$\rho(\mathsf{Peak}, \mathsf{HtM} - \mathsf{X})$	
Peak	1.00	(0.000)	0.78	(0.000)	NA	(NA)
HtM	0.78	(0.000)	1.00	(0.000)	NA	(NA)
Own	0.42	(0.086)	0.36	(0.146)	0.68	(0.003)
Mort	-0.35	(0.155)	-0.32	(0.196)	0.71	(0.001)
HO	0.54	(0.022)	0.47	(0.047)	0.60	(0.011)
Wealthy	0.35	(0.148)	0.23	(0.35)	0.72	(0.001)
Flex	-0.04	(0.894)	-0.03	(0.914)	0.79	(0.000)
Age	-0.05	(0.851)	0.10	(0.703)	0.79	(0.000)
LTV	-0.34	(0.163)	-0.26	(0.296)	0.72	(0.001)
LTI	-0.37	(0.135)	-0.22	(0.37)	0.72	(0.001)
Trade	0.10	(0.67)	-0.18	(0.483)	0.81	(0.000)
ROL	-0.08	(0.797)	0.11	(0.726)	0.88	(0.000)
Services	-0.41	(0.083)	-0.19	(0.442)	0.73	(0.001)
HP Growth	0.34	(0.166)	-0.13	(0.623)	0.83	(0.000)
GDPpc	-0.44	(0.058)	-0.47	(0.048)	0.68	(0.003)

Note: The *first column* shows the correlation coefficient between estimated peak values and the variables that vary across the rows in the table. The *second column* shows the correlation coefficient between HtM shares and the variables that vary across the rows in the table. The *third column* shows the semipartial correlation between the estimated peak values and HtM shares. The p-values for the correlation coefficients are reported within parentheses to the right of each coefficient. Calculations of *Peak, HtM, Own, Mort, HO, Wealthy, Flex, Age, LTV* and *LTI* are based on data from the HFCS. *HP Growth* is the average quarterly year-on-year growth in Eurostat's house price index from the first data point (2006Q1 for most countries) until 2012Q4. Greece is missing from the index. GDP per Capita is for 2008. See text for information about the source of the other variables.

semipartial correlations are reported in the third column of Table 4.2 and indicate the correlation coefficient between the peak effects and HtM shares, after the variation in HtM shares explained by other variables, varying across the rows in the table, has been accounted for. Going down the rows, we conclude that the coefficient remains large. From being 0.78 without having "controlled for" any other variable, it reaches its lowest value at 0.60 when we account for the variation in HtM explained by HO and as high as 0.88 when we instead extract the variation in HtM explained by *ROL*. Based on the results presented in Table 4.2, we find no variable that supports the conclusion that correlation between the peak effects and HtM shares is driven by omitted variables. For example, it could have been the case that all HtM households, but no non-HtM households, had mortgages. In such a case, the correlation between output responses and shares of HtM could potentially be explained by the fact that higher shares of households with mortgages caused larger output responses. The results presented in Table 4.2 suggest that a higher fraction of constrained households causes output to respond more to monetary policy shocks.

Intuitively, many of the variables considered in the table, such as home ownership (Own) and mortgage holdings (Mort), seem closely related to the HtM status of a household. Hence it may be surprising that none of the variables in the table are able to significantly attenuate the correlation we find. It is important to realize, however, that none of the variables in table 4.2, except for the constructed variable HtM itself, take the liquidity of a household's asset positions into account. In particular, the latter quantifies the relationship *liquid assets-to-income*. Together with the estimated effects for monetary policy shocks on output, the results in table 4.2 suggest that, if one is to construct a statistic based on household asset data, with the intent to capture MPC, then no single variable by itself is satisfactory but one must classify assets based on liquidity and set them in relation to income.

4.5 Conclusion

The introduction of heterogeneous agents into New Keynesian models is becoming widespread. However, there is still a lack of empirical evidence on how household heterogeneity in income and wealth affects the response of aggregate output following a monetary policy shock. In this paper we provide such evidence, showing that aggregate output responses are larger in countries with a higher share of liquidity constrained households.

We estimate country specific output responses in the euro area, following an expansionary monetary policy shock. The IRFs are produced using Local Projections (Jordà, 2005). To identify surprise changes in the policy rate, we construct an instrument based on movements in Eonia OIS rates during a narrow time window around the ECB's monetary policy announcement and the subsequent press conference. Given that the countries within the euro area share a central bank, we can rule out that any heterogeneity in IRFs is due to differences in the success of our identification method across countries.

We find that output responses to a common monetary policy shock in the euro area are heterogeneous across countries in terms of cumulative impact and peak values.

Subsequently, we correlate the country specific responses with proxies for the share of liquidity constrained households across countries. Intuitively, these households are less able to smooth income fluctuations following monetary policy shocks. Our main measure is the share of households that are classified as HtM, according to the definition by Kaplan et al. (2014), but we construct four additional measures of the share of constrained households, which are distinct in the surveys and time periods used to construct them.

On average, countries with a higher share of liquidity constrained households react more strongly to a monetary policy shock. When splitting the sample by shares of HtM households, the aggregate response of the high-HtM countries is significantly stronger than that of the low-HtM countries. These findings are in line with theoretical work, given plausible assumptions about the elasticity of constrained households' incomes to aggregate income (Bilbiie, 2020).

Our findings support the notion that research on monetary policy needs to account for heterogeneity across the income and wealth distributions. Furthermore, they imply that liquidity is an important factor in how monetary policy shocks affect households and the real economy. Additional empirical research is needed, however, to understand the mechanism through which this heterogeneity in liquidity directly shapes the responses of output to monetary policy shocks. We consider this a fruitful avenue for future research.

References

Altavilla, C., Giannone, D., and Lenza, M. (2016). The financial and macroeconomic effects of the omt announcements. *International Journal of Central Banking.*

Ampudia, M. and Van den Heuvel, S. (2018). Monetary policy and bank equity values in a time of low interest rates. *ECB Working Paper.*

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–67.

Bilbiie, F. O. (2008). Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of economic theory*, 140(1):162–196.

Bilbiie, F. O. (2019). Monetary Policy and Heterogeneity: An Analytical Framework. *Mimeo.*

Bilbiie, F. O. (2020). The new keynesian cross. *Journal of Monetary Economics*, 114:90–108.

Botero, J. C., Djankov, S., Porta, R. L., Lopez-de Silanes, F., and Shleifer, A. (2004). The regulation of labor. *The Quarterly Journal of Economics*, 119(4):1339–1382.

Burriel, P. and Galesi, A. (2018). Uncovering the heterogeneous effects of ECB unconventional monetary policies across euro area countries. *European Economic Review*, 101:210–229.

Calza, A., Monacelli, T., and Stracca, L. (2013). Housing finance and monetary policy. *Journal of the European Economic Association*, 11(suppl_1):101–122.

Carroll, C., Slacalek, J., Tokuoka, K., and White, M. N. (2017). The distribution of wealth and the marginal propensity to consume. *Quantitative Economics*, 8(3):977–1020.

Chow, C. and Lin, A. (1971). Best Linear Unbiased Interpolation , Distribution , and Extrapolation of Time Series by Related Series. *The Review of Economics and Statistics*, 53(4):372–375.

Cloyne, J., Ferreira, C., and Surico, P. (2020). Monetary policy when households have debt: new evidence on the transmission mechanism. *The Review of Economic Studies*, 87(1):102–129.

Cloyne, J., Huber, K., Ilzetzki, E., and Kleven, H. (2019). The effect of house prices on household borrowing: A new approach. *American Economic Review*, 109(6):2104–36.

Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). Innocent Bystanders? Monetary policy and inequality.

Corsetti, G., Duarte, J. B., and Mann, S. (2018). One money, many markets-a factor model approach to monetary policy in the euro area with high-frequency identification. *Mimeo*.

Datascope (2018). Euro Overnight Indexed Swap (3-month and 1-year). `https://hosted.datascope.reuters.com/datascope/` (accessed July 31, 2018).

European Central Bank (2014). Eurosystem Household Finance and Consumption Survey. `https://www.ecb.europa.eu/stats/ecb_surveys/hfcs/html/index.en.html` (waves 1 and 2 accessed September 30, 2017, wave 3 accessed October 22, 2020).

Eurostat (2019a). All Item HICP (2015 = 100) - monthly data (index). `https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=prc_hicp_midx&lang=en` (accessed February 5, 2019).

Eurostat (2019b). Arrears on utility bills - EU-SILC survey, Total . `https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=ilc_mdes07&lang=en` (accessed February 5, 2019).

Eurostat (2019c). Gross domestic product at market prices, Chain linked volumes (2010), euro per capita, 2008. `https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=sdg_08_10&lang=en` (accessed February 5, 2019).

Eurostat (2019d). Gross domestic product at market prices, Chain linked volumes (2010), million euro, Seasonally and calendar adjusted. `https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=namq_10_gdp&lang=en` (accessed February 5, 2019).

Eurostat (2019e). Inability to face unexpected financial expenses - EU-SILC survey, Total. `https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=ilc_mdes04&lang=en` (accessed February 5, 2019).

Eurostat (2019f). Money market interest rates - monthly data, Euro Area. `https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=irt_st_m&lang=en` (accessed February 5, 2019).

Eurostat (2019g). Production in Industry, Mining and quarrying; manufacturing; electricity, gas, steam and air conditioning supply, Volume index of production (2010 = 100), Seasonally and calendar adjusted. `https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=sts_inpr_m&lang=en` (accessed February 5, 2019).

Eurostat (2019h). Retail trade, except of moter vehicles and motorcycles, Index of deflated turnover (2010 = 100), Seasonally and calendar adjusted. https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=sts_trtu_m&lang=en (accessed February 5, 2019).

Eurostat (2019i). Unemployment - monthly data, Percentage of active population from 25 to 74 years, Seasonally adjusted, not calendar adjusted. https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=une_rt_m&lang=en (accessed February 5, 2019).

Eurostat (2021). Final consumption expenditure of Households, Chain linked volumnes (2010), million euro, Seasonally and calendar adjusted. https://appsso.eurostat.ec.europa.eu/nui/show.do?dataset=namq_10_fcs&lang=en (accessed May 1, 2021).

Fagereng, A., Holm, M. B., and Natvik, G. J. J. (2020). MPC Heterogeneity and Household Balance Sheets. *American Economic Journal: Macroeconomics*.

Flodén, M., Kilström, M., Sigurdsson, J., and Vestman, R. (2020). Household Debt and Monetary Policy: Revealing the Cash-Flow Channel. *The Economic Journal*, 131(636):1742–1771.

Galí, J., López-Salido, J. D., and Vallés, J. (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association*, 5(1):227–270.

Galí, J. and Monacelli, T. (2008). Optimal monetary and fiscal policy in a currency union. *Journal of International Economics*, 76(1):116–132.

Georgiadis, G. (2015). Examining asymmetries in the transmission of monetary policy in the euro area: Evidence from a mixed cross-section global VAR model. *European Economic Review*, 75(January):195–215.

Gerlach, S. and Smets, F. (1995). The Monetary Transmission Mechanism: Evidence from the G-7 Countries.

Gertler, M. and Karadi, P. (2015). Monetary Policy Surprises, Credit Costs,and Economic Activity. *American Economic Journal: Macroeconomics*, 7(1):44–76.

Gornemann, N., Kuester, K., and Nakajima, M. (2016). Doves for the rich, hawks for the poor? distributional consequences of monetary policy.

Guvenen, F., Ozkan, S., and Song, J. (2014). The nature of countercyclical income risk. *Journal of Political Economy*, 122(3):621–660.

Hagedorn, M., Luo, J., Manovskii, I., and Mitman, K. (2019). Forward guidance. *Journal of Monetary Economics*.

Honkkila, J. and Kavonius, I. K. (2013). Micro and Macro Analysis on Household Income, Wealth and Saving in the Euro Area. *ECB Working Paper Series*, 1619(1619).

Jarociński, M. and Karadi, P. (2020). Deconstructing monetary policy surprises-The role of information shocks. *American Economic Journal: Macroeconomics*.

Jordà, Ò. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

Kaplan, G., Moll, B., and Violante, G. L. (2018). Monetary policy according to hank. *American Economic Review*, 108(3):697–743.

Kaplan, G. and Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4):1199–1239.

Kaplan, G., Violante, G. L., and Weidner, J. (2014). The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity*, 2014(1):77–138.

Kuttner, K. N. (2001). Monetary policy surprises and interest rates: Evidence from the Fed funds futures market. *Journal of Monetary Economics*, 47(3):523–544.

Lloyd, S. (2018). Estimating nominal interest rate expectations: Overnight indexed swaps and the term structure. *Bank of England Working Paper*.

Mandler, M., Scharnagl, M., and Volz, U. (2016). Heterogeneity in euro-area monetray policy transmission: Results from a large multi-country BVAR model. *Deutsche Bundesbank Discussion Paper*, (03).

Nakamura, E. and Steinsson, J. (2018). High-frequency identification of monetary non-neutrality: The information effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.

Organisation for Economic Co-operation and Development (2019). Unemployment rate (monthly), Aged 25 and over, All persons, Jan 2000, Estonia . `https://stats.oecd.org/BrandedView.aspx?oecd_bv_id=mei-data-en&doi=data-00046-en#` (accessed February 5, 2019).

Patterson, C. (2019). The matching multiplier and the amplification of recessions. *Mimeo*.

Ramey, V. A. (2016). Macroeconomic Shocks and Their Propagation. *Handbook of Macroeconomics*, 2:71–162.

Stock, J. H. and Watson, M. W. (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal*, 128(610):917–948.

Timmer, M. P., Dietzenbacher, E., Los, B., Stehrer, R., and de Vries, G. J. (2015). An illustrated user guide to the world input–output database: the case of global automotive production. *Review of International Economics*, 23(3):575–605.

Werning, I. (2015). Incomplete Markets and Aggregate Demand. *National Bureau of Economic Research Working Paper Series*.

Wong, A. (2019). Refinancing and the transmission of monetary policy to consumption. *Mimeo*.

World Bank (2019). Services, Value added (Percent of GDP) Series name: NV.SRV.TOTL.ZS. 2000-2012. `https://databank.worldbank.org/source/world-development-indicators` (accessed January 30, 2019).

World Input Output Database (2018). World Input-Output Tables. 2000-2012. `http://www.wiod.org/database/wiots16` (accessed June 5, 2018).

Appendices

4.A Income Elasticities

The amplification result outlined in Bilbiie (2019) requires that constrained (unconstrained) households' income elasticities with respect to aggregate income are larger (smaller) than one. Empirical evidence to this effect is scarce.[39] We therefore test for the income elasticity mechanism using the HFCS dataset.

A subset of households in our sample, from a subset of countries that participate in the HFCS, are interviewed in multiple survey waves. We use data for these households and investigate their income elasticities with respect to aggregate income. Since data from three waves currently exist, we compute the individual growth rates between (i) the first and second waves and (ii) second and third waves. To limit the influence of outliers, households whose income or income growth rates were below or above the 1st and 99th percentiles, respectively, in each country and time period, were removed.[40] Since the HtM status of a household can change between the survey waves, we choose to classify a household as HtM if it was classified thusly in the first wave contributing to the income growth rate.[41] Sample

[39] Coibion et al. (2017) find that inequality rises after contractionary monetary policy in the US. They estimate that the change in labor earnings of high net–worth households is lower than that of low net–worth households after monetary shocks, and that incomes of households at the 90th percentile rise somewhat relative to the median household, while households at the 10th percentile see their relative incomes fall particularly sharply. Patterson (2019) documents a positive covariance between workers' MPCs and their earnings elasticity to GDP that is large enough to increase shock amplification.

[40] The result presented in Equation (4.8) is robust to trimming below and above the 5th and 95th percentiles, respectively.

[41] As is discussed more in detail in Appendix 4.F, the HFCS imputes data for missing values for some variables and this is done five times, which results in

weights are employed in the estimation.

We run the following regression, following, e.g., Guvenen et al. (2014), but distinguishing by HtM status:

$$\begin{aligned}\Delta y_{i,n,t} = \underset{\substack{[9.307]\\(1.143)}}{\alpha} + \underset{\substack{[-1.401]\\(2.670)}}{\beta} \mathrm{HtM}_{i,n,t-1} + \underset{\substack{[1.168]\\(0.137)}}{\gamma} \Delta Y_{n,t} \\ + \underset{\substack{[0.643]\\(0.326)}}{\delta} \Delta Y_{n,t} \times \mathrm{HtM}_{i,n,t-1} + e_{i,n,t} \qquad (4.8)\end{aligned}$$

where the left-hand-side variable is the growth rate of labor income for household i in country n between two periods $t-1$ and t, HtM is the variable that indicates the Hand-to-Mouth status of the household (in period $t-1$) and $\Delta Y_{n,t}$ is the growth rate of aggregate income in country n between periods $t-1$ and t. Lastly, the regression includes an interaction between aggregate income growth and Hand-to-Mouth status. The coefficients of interest are γ and δ, where γ captures the (average) elasticity of individual income growth with respect to aggregate income growth for unconstrained households, and $\gamma+\delta$ captures the (average) elasticity of individual income growth with respect to average income growth for financially constrained households.

The estimated coefficients are reported below their respective parameters and standard errors are placed inside parentheses.[42] The first

five implicates. As a result of the imputation, the HtM status that we assign to households can possible vary across implicates. For the exercise that we perform in the current section, we classify a household as HtM if it was classified as HtM in at least three out of five implicates.

[42]Robust standard errors are clustered at the individual level. We explore other alternatives, like country level clustered standard errors or estimate the standard errors using a wild bootstrap with standard errors clustered at the country level. The former yields a δ coefficient that is statistically significant at the 95% confidence level, with only 12 clusters, and the δ coefficient is statistically significant at the 90% level in the latter case.

coefficient of interest, γ, is estimated to be 1.17 and is statistically significant at the 95 percent level.[43] On the other hand, δ is estimated to be 0.64 and is statistically significant at the 95 percent level (p-value 0.049). The value indicates that a one-percentage point increase in aggregate income is associated with financially constrained households' incomes increasing by 0.64 percentage points *more* than for unconstrained agents. Taken together, these findings suggest that if aggregate income grows, the income of financially constrained households grows by *more* and would, through the lens of Bilbiie (2019), lead to amplification, as our results in Section 4.4.1 suggest.

4.B Additional Figures and Tables

[43] Within each country, higher levels of income are associated with lower levels of income growth. It has the consequence that average income growth exceeds aggregate income growth, which explains why the estimated value for γ is greater than one.

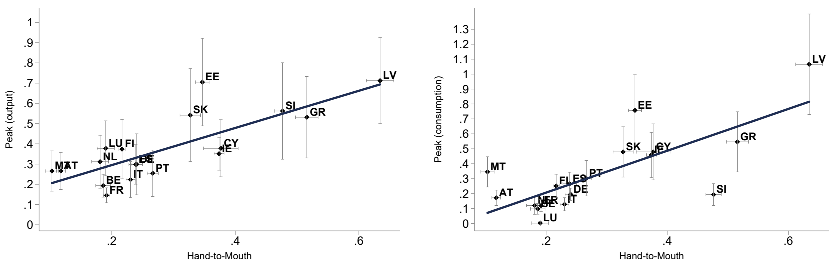

(a) ***(b)***

Figure 4.10: *Monetary policy effectiveness and Hand-to-Mouth shares – output and consumption responses with confidence bands*

Note: This figure plots Hand-to-Mouth shares against peak responses of output (panel (a)) and consumption (panel (b)). The HtM shares are calculated using data from the Eurosystem Household Finance and Consumption Survey. The vertical lines and horizontal lines represent (1 std) confidence bands for the peak responses and HtM shares, respectively. See appendix 4.F for more information about the standard errors for HtM.

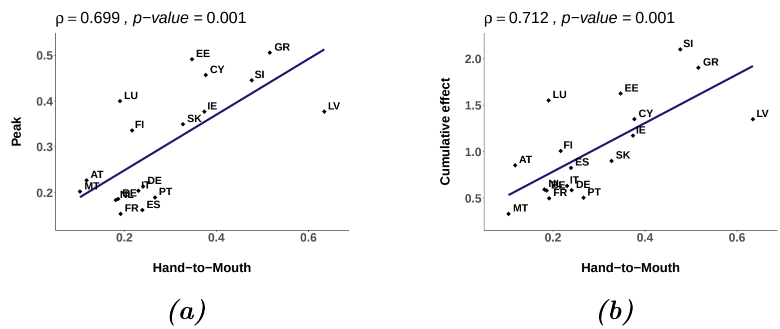

Figure 4.11: *Robustness including country specific lags*

Note: This figure plots the effectiveness of monetary policy, as measured by the peak effect and cumulative effect of the real GDP impulse responses, calculated using the LPIV estimation with country specific lags, against the share of households classified as living HtM in each euro area country (except Lithuania, not included in the HFCS). The LPIV estimation includes three country specific lags. The HtM shares are calculated using data from the Eurosystem Household Finance and Consumption Survey. The impulse is an expansionary monetary policy shock of one standard deviation. The blue lines are fitted from regressions of Peak/Cumulative values on HtM shares. In the upper left corner of each panel, we report the correlation coefficient ρ and the p-value. *Panel (a)*: Peak effects and share of Hand-to-Mouth. *Panel (b)*: Cumulative effects and share of HtM, normalized by aggregate euro area cumulative effect.

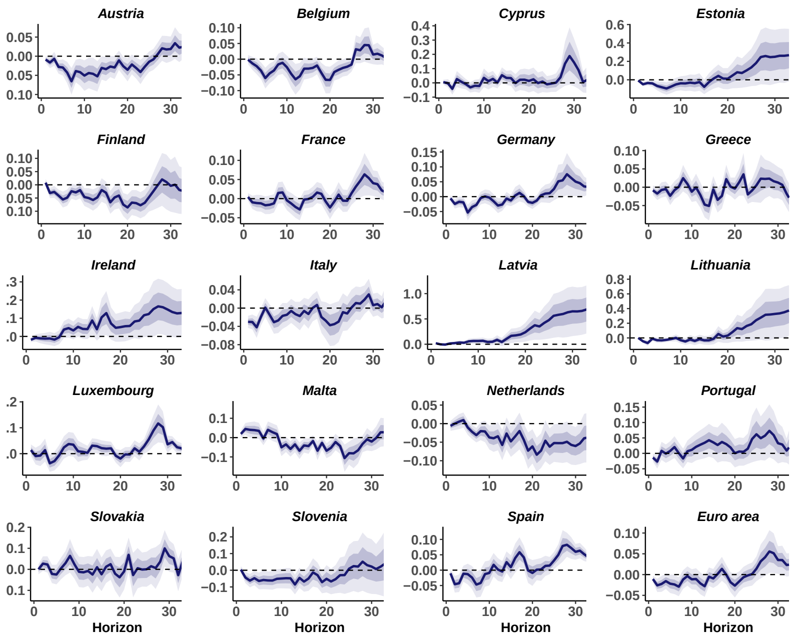

Figure 4.12: *Impulse responses for prices in euro area countries – LPIV*

Note: This figure shows impulse responses of prices to an expansionary monetary policy shock of one standard deviation. For each euro area country, the response is estimated using LPIV (Equation 4.3). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark and light blue shaded areas represent 1 and 2 standard deviation confidence bands, constructed using Newey-West estimators. Note that the y-axes are scaled differently across countries.

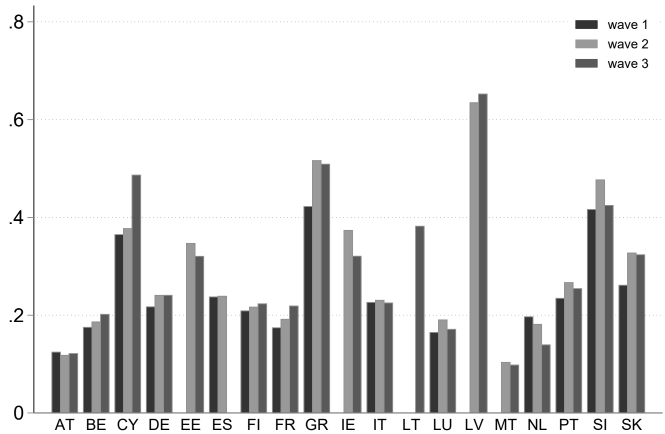

Figure 4.13: *Fraction of HtM households across HFCS survey waves*

Note: This figure shows the fraction of HtM households, calculated according to the approach in Kaplan et al. (2014), utilizing data from three different survey waves of the Eurosystem Household Finance and Consumption Survey.

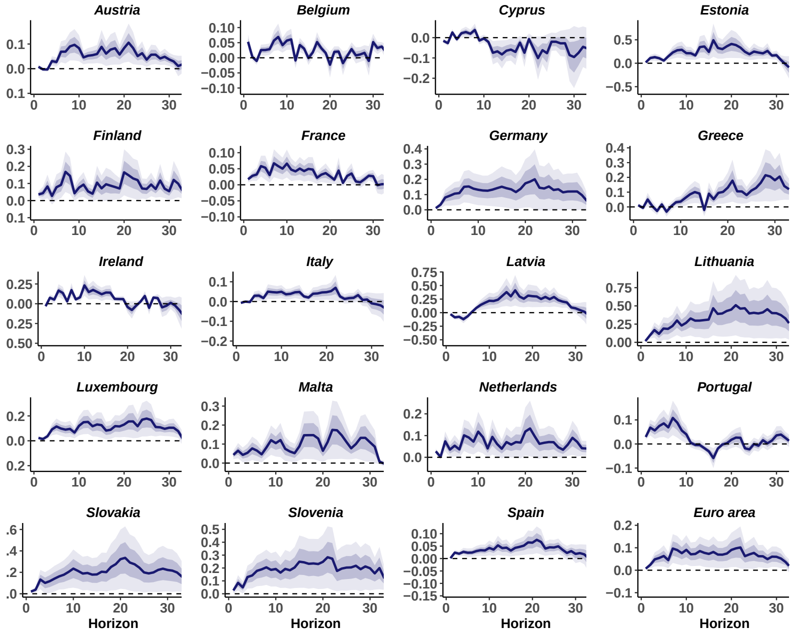

Figure 4.14: *Impulse responses for output in euro area countries – LPIV – JK*

Note: This figure shows impulse responses of real GDP to an expansionary monetary policy shock of one standard deviation. The shock series is the Monetary Policy shock series reported in Jarociński and Karadi (2020). For each euro area country, the response is estimated using LPIV (Equation 4.3). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark and light blue shaded areas represent 1 and 2 standard deviation confidence bands, constructed using Newey-West estimators. Note that the y-axes are scaled differently across countries.

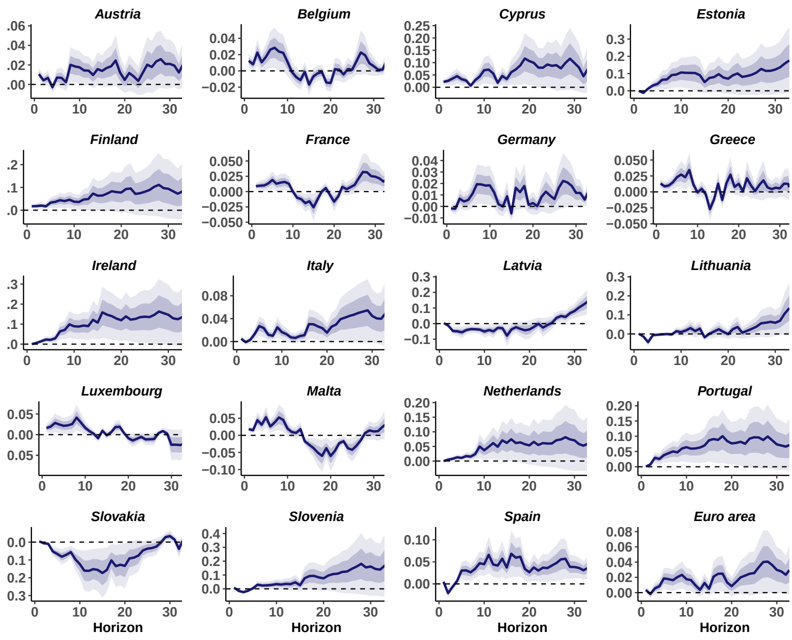

Figure 4.15: *Impulse responses for prices in euro area countries – LPIV – JK*

Note: This figure shows impulse responses of prices to an expansionary monetary policy shock of one standard deviation. The shock series is the Monetary Policy shock series reported in Jarociński and Karadi (2020). For each euro area country, the response is estimated using LPIV (Equation 4.3). The solid blue lines represent the IRFs produced by our preferred specification (see text for details). The dark and light blue shaded areas represent 1 and 2 standard deviation confidence bands, constructed using Newey-West estimators. Note that the y-axes are scaled differently across countries.

4.C The Global VAR Setting

As a robustness check to our main empirical framework, we construct an instrumented GVAR. We build a more structural—and restricted—setting than the LPIV, more similar to the widespread VAR estimation in the literature. We follow the GVAR setting in Burriel and Galesi (2018), except that we remove contemporaneous variables on the right hand side for endogeneity issues. All N economies are represented by the following system:

$$\Lambda Q_t = \kappa_0 + \sum_{j=1}^{r} K_j Q_{t-j} + \nu_t \tag{4.9}$$

where $Q_t = (y_{1t}, \pi_{1t}, ..., y_{Nt}, \pi_{Nt}, i_t)'$ is a $(2N+1) \times 1$ vector containing output and inflation for each country, and the global interest rate. Pre-multiplying both sides by Λ^{-1} yields

$$Q_t = h_0 + \sum_{j=1}^{r} H_j Q_{t-j} + \nu_t \tag{4.10}$$

where $h_0 = \Lambda^{-1}\kappa_0$, $H_j = \Lambda^{-1}K_j$ and $\nu_t = \Lambda^{-1}\nu_t$. We seek to estimate (4.10). Unfortunately, this is infeasible due to the curse of dimensionality: there are too many parameters to estimate for the restricted number of observations that we have. In order to overcome this situation, we borrow two key assumptions from the GVAR literature: we assume that (i) foreign variables affecting country i will be a composite of an aggregate coefficient and the trade weight to each foreign economy, and (ii) that the ECB reacts to euro area aggregates and not to individual countries. In this way, our setting is akin to a standard GVAR, but without assuming the Small Open Economy framework that is necessary to rule out potential endogeneity bias.

We now explore each equation inside the (4.10) system. We start with the first block, that includes the Dynamic IS curve and the New Keynesian Phillips curve. Each domestic economy is represented by the following reduced-form VAR:

$$Y_{it} = c_i + \sum_{j=1}^{p_i} A_{ij} Y_{i,t-j} + \sum_{j=1}^{q_i} B_{ij} Y^*_{i,t-j} + \sum_{j=1}^{q_i} C_{ij} X_{t-j} + u_{it} \quad (4.11)$$

where c_i is a country specific intercept vector, Y_{it} is a 2×1 vector of domestic variables (i.e., output and inflation), Y^*_{it} is a 2×1 vector of aggregate foreign variables, X_t is a the ECB policy rate and u_{it} is a vector of idiosyncratic country-specific reduced form shocks. The foreign variables are computed as trade weighted aggregates $Y^*_{it} = \sum_{j \neq i} w_{ij} Y_{jt}$ with $\sum_{j \neq i} w_{ij} = 1$, where we assume that weights w_{ij} calculated using bilateral trade flows from the World Input Output Database (WIOD), are fixed over time.[44]

Stacking all countries in our model, using that $Y^*_{it} = W_i Y_t$, with W_i being country-specific weight matrices, we can write equation (4.11) as

$$Y_t = c + \sum_{j=1}^{p} G_j Y_{t-j} + \sum_{j=1}^{q} C_j X_{t-j} + u_t \quad (4.12)$$

where $G_j = \left(A_j + B_j W\right)$, $Y_t = (Y'_{1t}, \ldots, Y'_{Nt})'$, $u_t = (u'_{1t}, \ldots, u'_{Nt})'$, $c = (c'_1, \ldots, c'_N)'$, $C_j = (C'_{1j}, \ldots, C'_{Nj})'$, $p = \max(p_i, q_i)$ and $q = \max(q_i)$.

Next, the second building block consists of variables which affect

[44]The weights are calculated using bilateral trade flows for years 2002 through 2012. See Timmer et al. (2015) for a user guide to the World Input Output Database (2018).

all countries, i.e., the interest rate controlled by the ECB,

$$X_t = c_x + \sum_{j=1}^{p_x} D_j X_{t-j} + \sum_{j=1}^{q_x} F_j \widetilde{Y}_{t-j} + u_{xt} \tag{4.13}$$

where u_{xt} is a vector of idiosyncratic reduced-form shocks and $\widetilde{Y}_t$ is a weighted average of all countries' domestic variables, with weights based on GDP shares $\widetilde{Y}_t = \widetilde{W} Y_t = \sum_j \widetilde{w}_j Y_{jt}$ with $\sum_j \widetilde{w}_j = 1$.

Notice that equation (4.13) is no more than a standard Taylor rule that the ECB is assumed to follow: the current interest rate depends on lags of output and inflation, plus lags on the interest rate itself. Stacking the two blocks given by (4.12) and (4.13), we obtain the following system of equations, which is exactly the same as in (4.10),

$$Q_t = h_0 + \sum_{j=1}^{r} H_j Q_{t-j} + \nu_t \tag{4.14}$$

where $r = \max(p, s)$, and the vector $Q_t = (Y_t', X_t')'$ includes all country-specific and common variables, $h_0 = \begin{bmatrix} c \\ c_x \end{bmatrix}$, $H_j = \begin{bmatrix} G_j & C_j \\ F_j \widetilde{W} & D_j \end{bmatrix}$ and $\nu_t = \begin{bmatrix} u_t \\ u_{xt} \end{bmatrix}$. In our baseline estimation, we set $p_i = q_i = 3$ $\forall i \in N$, and $p_x = q_x = 3$.

A novelty in this paper is that we identify monetary responses in a GVAR setting using exogenous instruments. In particular, we identify the structural monetary policy shock from the reduced-form errors. The structural error vector can be written as $\nu_t = \begin{pmatrix} u_t \\ u_{xt} \end{pmatrix} = \Lambda^{-1} \begin{pmatrix} \varepsilon_t \\ \varepsilon_{xt} \end{pmatrix}$. Λ^{-1} being unknown, we would not be able to obtain the true impulse responses. We use external instruments to identify (part of) Λ^{-1}. Since we are only interested in a monetary policy shock, we need to identify the relevant column of the variance-covariance matrix that describes the effect of ε_{xt} on the other structural errors in ν_t.

The first part of the identification strategy is similar to the LPIV: we estimate the model in equations (4.11) and (4.13) using OLS. As before, one can verify that the reduced form errors ν_t are linear combinations of the structural errors ε_{it} $\forall i \in N$ and ε_{xt}, where Λ^{-1} is a $2N+1$ square matrix with elements on its 2×2 block diagonal and zeroes elsewhere. Without further restrictions, we cannot identify the full matrix Λ^{-1} describing the relationship between reduced form and structural errors. We can, however, identify the column of the matrix describing the influence of the structural component of the interest rate ε_{xt} on the other variables. The relevant column of Λ^{-1} can be identified by introducing the contemporaneous interest rate on the RHS of the system of equations (4.11), making use of 2SLS. Following Stock and Watson (2018), we identify the relative response of variable j to a structural shock in x in two steps. First, we instrument X_t using a valid instrument satisfying $\mathbb{E}\left[Z_t \varepsilon_{xt}\right] = \alpha$ and $\mathbb{E}\left[Z_t \varepsilon_{jt}\right] = 0$ where $j \neq x$, and regress the contemporaneous interest rate on the instrument Z_t, lags of the instrument and the rest of the variables that will enter the second stage of the 2SLS estimation:

$$X_t = c_i + \sum_{j=1}^{p_i} A_{ij} Y_{i,t-j} + \sum_{j=1}^{q_i} B_{ij} Y^*_{i,t-j} + \sum_{j=1}^{s_i} C_{ij} X_{t-j} + \theta^{SW}_{ix} Z_t + u_{it} \tag{4.15}$$

From this first stage, we obtain the fitted policy rate $\widehat{X}_{it}$ and we can then estimate the system (4.16). Second, we estimate the following system of equations for every country i,

$$Y_{it} = c_i + \sum_{j=1}^{p_i} A_{ij} Y_{i,t-j} + \sum_{j=1}^{q_i} B_{ij} Y^*_{i,t-j} + \sum_{j=1}^{s_i} C_{ij} X_{t-j} + \Theta^{SW}_{ix} \widehat{X}_{it} + u_{it} \tag{4.16}$$

The contemporaneous effect of a monetary policy shock on other variables is captured through Θ^{SW}_{ix}, which is used together with the endogenous variables' coefficient matrix to obtain the impulse responses.

4.D Panel LPIV

In Figure 4.9b, we compare the average impulse responses of output in two sets of countries: those with high and low levels of liquidity constrained individuals, according to our HtM variable. This approach does not allow for country-specific heterogeneity beyond the two HtM categories. Hence, in this section, we estimate a Panel LPIV which allows us to control for country fixed effects in addition to the high/low-HtM dummy.

We run the following regression, following Jordà (2005), as before:

$$\begin{aligned} y_{n,t+h} - y_{n,t-1} = &\alpha^h + \beta^h \hat{i}_t + \delta^h \widehat{i_t \times htm_n} + \gamma^h_n \\ &+ \xi^h_n htm_n + \sum_{j=1}^{p} \Gamma^h_{n,j} htm_n X_{t-j} + u_{n,t+h}, \\ &h = 0, \dots, H \end{aligned} \tag{4.17}$$

where y_n is log of output in country n, $\hat{i}$ and $\widehat{i_t \times htm_n}$ are the fitted values from the first-stage regression, htm_n takes value of 0 if the HtM share in a country is below the median value across all countries and 1 otherwise, and γ^h_n represents the country fixed effects. The control variables X_{t-j} are the same for all countries, namely lags of euro area real GDP, euro area HICP, lags of the policy variable and lags of the instrument Z. We interact these control variables with the htm dummy. We construct the instrument for $i_t \times htm_n$ by multiplying our instrument for the policy rate, Z_t with the high-HtM dummy: $Z_t \times htm_n$.

The coefficient of interest in this estimation is δ_n^h, which measures the additional impact of an interest rate change on real GDP in countries with higher-than-median shares of HtM individuals, beyond the impact already captured in β_n^h.

Figure 4.16 plots both coefficients across horizons h. The left panel indicates that GDP falls for all countries, in response to a one standard deviation shock. As already suggested in Figure 4.9b in the main body of the paper, however, GDP falls by significantly more in countries with a higher share of HtM households. The crucial difference between the two exercises is that here, we are able to control for country-specific fixed-effects beyond the high/low-HtM classification. We view the fact that the conclusions are unchanged as encouraging.

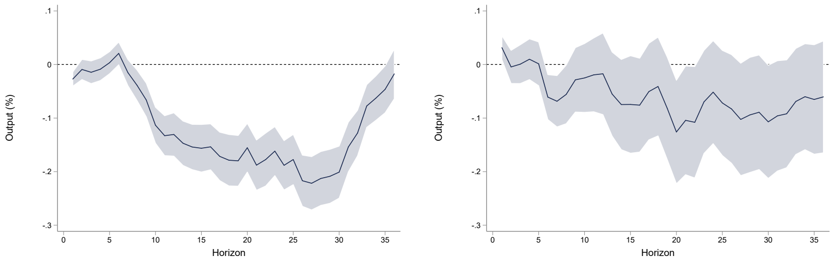

(a) ***(b)***

Figure 4.16: *Panel LPIV*

Note: The *Left Panel* plots the coefficient β^h from Equation (4.17) for each horizon h in response to a one standard deviation shock to our instrument. The *Right Panel* plots the coefficient δ^h from Equation (4.17). The blue shaded area represents 1 standard deviation confidence bands.

4.E European Overnight Indexed Swap Data

We obtain a minute-frequency series for Eonia Overnight Indexed Swaps from Datascope. We compute the fixed rate of the swap as the mid point between the bid and ask price at the close of each minute. We then drop all dates from the sample that are not ECB announcement dates.

The resulting series contains implausible outliers, e.g., the rate decreasing to zero for one minute, or short fluctuations of more than 5 standard deviations. Consequently, we drop the highest and lowest percentile of observations on each announcement day. Lastly, we manually drop remaining implausible observations if they fall within either of the two announcement windows.

For our final series, we exclude the observation on November 6th, 2008. On this day, the ECB cut interest rates by 50 BP, one of the largest cuts during our sample period. However, the market reaction in the overnight indexed swap rates indicates that markets perceived it as contractionary. Likely, this is due to the Bank of England having lowered its policy rate by 150 BP hours prior. Including the observation does not change our results or the conclusions, except for the first stage F-statistic, which falls to 4.4.

4.F Obtaining HtM Shares using Data from the HFCS

The HFCS imputes data for missing values related to assets, liabilities and income variables. Our calculations are partly based on these imputed data. A missing value is imputed five times (multiple imputation), where each time a different random term is added to the predicted value. If this was not be done, imputation uncertainty would not be taken into account. This has the consequence that statistics can vary between implicates.

To find point estimates for the statistics based on HFCS data, we average over all the implicates. We consistently use the cross-sectional (full sample) weights, which are mainly intended to compensate for some households being more likely to be selected into the sample than others. In other words, if a type of household has been over-sampled, then they are given less weight in the estimation.

We use techniques that are standard when computing variance estimates for multiple imputed survey data. In short, there are two sources of uncertainty that we need to account for. The first (B) is the uncertainty that is associated with the imputation. This is given by the variance of the point estimates (using the full sample weights). The second (W) is the uncertainty associatied with sampling and the weights that should be given each observation. The HFCS contains

$1,000$ replicate weights and the uncertainty for a statistic associated with sampling and weights is given by the variance of the estimators from using different replicate weights, averaged across the implicates. The total variance, T, is given by $\mathsf{T} = W + \frac{6}{5}\mathsf{B}$. We refer the reader to the HFCS user manual for more details about finding the variance estimates.

Before we label households, we drop observations where the age of the reference person in the household is below 20 or above 80. As in Kaplan et al. (2014) we drop observations when the only income that the household receives is from self-employment. The results do not change markedly if we choose to keep these observations.

We need to categorize variables as liquid wealth, illiquid wealth, liquid debt and illiquid debt. We follow Kaplan et al. (2014) to a large extent. In Table B1 we present what variables go into respective category and the *Name* refers to its unique name in the HFCS data. The difference between how we categorize the variables and how Kaplan et al. (2014) do it is that we categorize savings accounts as liquid assets while they categorize it as illiquid for the European countries. We choose to categorize it as liquid as it is our view that households can, in general, make adjustments to the balance on saving accounts without incurring substantial costs. In the Panel Study of Income Dynamics (PSID), saving accounts are combined with other assets such as checking accounts. Moreover, in the calibration of the model in Carroll et al. (2017), saving accounts are categorized as liquid.

In the calculation of HtM shares, Kaplan et al. (2014) assume that households on average are paid bi-weekly. In our calculations we will assume that households on average are paid once every month, which we believe is a more accurate assumption about the payment frequency in European countries.

We define $\text{liquid wealth} = \text{liquid assets} - \text{liquid debt}$ and $\text{illiquid wealth} = \text{illiquid assets} - \text{illiquid debt}$.

Name	Description	Comment
Income		
di1100	employee income	
di1610	unemployment benefits	
di1620	other social transfers	
hg0210	income from regular private transfers	
di1510	gross income from public pensions	
Liquid assets		
hd1110	value of sight accounts	Scaled by 1.0556 to adjust for cash missing in the HFCS.[45]
da2102	mutual funds, total	
da2105	shares, publicly traded	
da2103	bonds	
hd1210	value of saving accounts	Illiquid in Kaplan et al. (2014)
Illiquid assets		
hb0900	current price of household main residence	
hb280x	other property $x: current value	x={1, 2, 3}
hb2900	additional properties current value	
sum of pf0710 across HH members	current value of all occupational pension plans that have an account	
da2109	voluntary pension/whole life insurance	
Liquid debt		
hc0220	amount of outstanding credit line/overdraft balance	
hc0320	amount of outstanding credit cards balance	
Illiquid debt		
hb170x	HMR mortgage $x: amount still owed	x = {1, 2, 3}
hb370x	other property mortgage $x: amount still owed	x = {1, 2, 3}
hb4100	money still owed on additional other property loans	
hb2100	money still owned on additional HMR loans	
hc080x	non-collateralised loan $x: outstanding balance of loan hc080x counted if hc050xa$\leqslant$ 2	x = {1, 2, 3}

Table 4.3: *Classification of income and assets in the HFCS*

4.G Tenure Status, Mortgages and HtM Status

4.G.1 Ownership rates and mortgages among W-HtM households

Figure 4.17 shows that the majority of households who have been classified as W-HtM households own the property in which they live (represented by the total length of each bar). However, in most countries the majority of W-HtM households do not have a mortgage (black). The countries are ordered according to their shares of W-HtM households, with the country with the lowest share of W-HtM households on top (Austria).

4.G.2 HtM status among homeowners

Cloyne et al. (2020) find that the consumption responses of homeowners are significantly smaller than the consumption responses of mortgagors and renters. They use data from the U.K. and U.S. and classify very few homeowners as Hand-To-Mouth (see Figure 10 in their paper). In our data, however, homeowners make up a substantial fraction of HtM households in many countries (see Figure 4.18). In some countries, it is even the case that a majority of HtM households are homeowners. Hence, we do not think that our results contradict the mentioned study, since homeowners appear to have different characteristics in the countries in our sample, compared to homeowners in the U.K. or the U.S.

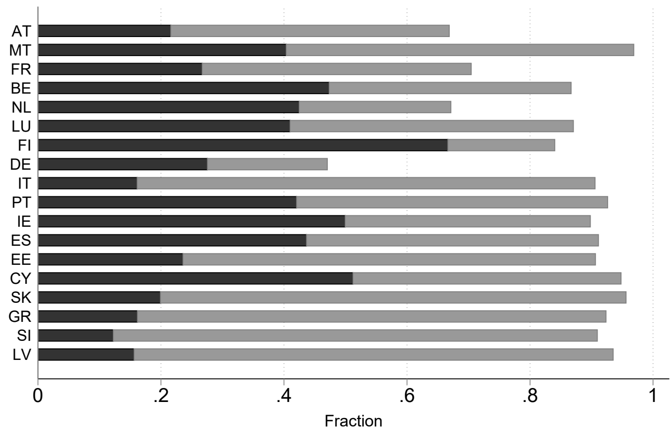

Figure 4.17: *Ownership and mortgages among the Wealthy HtM*

Note: This figure shows three things: (i) the fraction of W-HtM households who own the residence in which they live (total length of each bar), (ii) the fraction of W-HtM households who have a mortgage (black) and the fraction of W-HtM who own their residence but do not have a mortgage (gray). The countries are ordered according to their shares of W-HtM households, with the country with the lowest share of W-HtM households on top (AT). The fractions are computed using data from the Eurosystem Household Finance and Consumption Survey (HFCS).

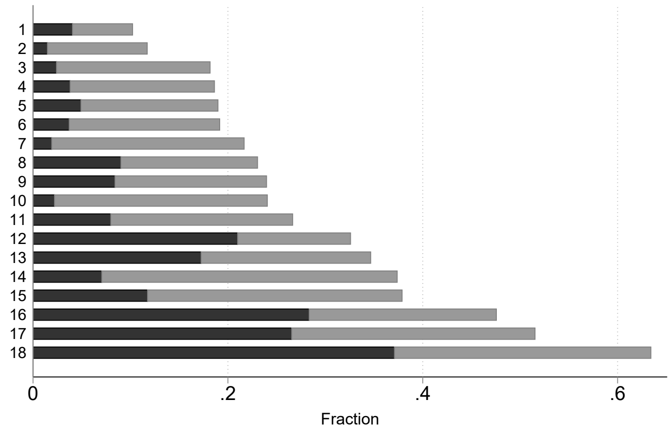

Figure 4.18: *HtM status among homeowners*

Note: This figure divides HtM shares (total length of each bar) up in to households who are homeowners (black) and not homeowners (renters or mortgagors, gray). The countries are ordered according to their share of HtM households, with the country with the lowest share of HtM households on top (MT). The fractions are computed using data from the Eurosystem Household Finance and Consumption Survey (HFCS).

4.H Local Projections Data

Inflation: We obtain the monthly Harmonized Index of Consumer Prices for all items for all countries in our sample and the euro area from Eurostat (*prc_hicp_midx*). See Eurostat (2019a).

Industrial Production: We obtain monthly values for Industrial Production (excluding construction) from Eurostat. The series is seasonally and calendar adjusted (`sts_inpr_m`). Because Ireland changed its formula for the calculation of some national aggregates, we make some assumptions to keep the series as coherent as possible. The change affects the value of Industrial Production in the first two months of 2015, resulting in growth rates in excess of 10%. We substitute these two growth rates with the average growth over 2014, which results in a level shift for all IP values after March 2015. See Eurostat (2019g).

Unemployment rate: We obtain monthly values of the unemployment rates for all countries in our sample from Eurostat (`une_rt_m`). The rates are measured for the active population aged 25 to 74 and are seasonally and calendar adjusted. For Estonia, the value of January 2000 is missing. We obtain it from the Organisation for Economic Co-operation and Development (2019) (`LRHUADTT`). The rest of the series coincides with the values from Eurostat. See Eurostat (2019i).

Real GDP: We obtain the quarterly values for Real GDP for all countries in our sample from Eurostat (`namq_10_gdp`). The series measures chain-linked volumes of Gross Domestic Product and is seasonally and calendar adjusted. Again, we adjust the series of Ireland due to implausibly high GDP growth in the first quarter of 2015. We substitute the reported growth rate in 2015Q1 with the average growth rate during 2014, which results in a level shift of all subsequent observations. See Eurostat (2019d).

Eonia: We obtain values for the European OverNight Index Average from Eurostat (irt_st_m). See Eurostat (2019f).

Retail trade: We obtain monthly data on Retail trade, except of motor vehicles and motorcycles from Eurostat for all countries in our sample. The series refers to deflated turnover and is seasonally and calendar adjusted (sts_trtu_m). See Eurostat (2019h).

Consumption: We obtain data on the final consumption expenditure of households from Eurostat (namq_10_fcs). The series is seasonally and calendar adjusted. See Eurostat (2021).

GDP per Capita: We obtain data on Real GDP per capita in 2008 from Eurostat (SDG_08_10). See Eurostat (2019c).

Sammanfattning

Allokeringen av utgifter och tid över tid

I detta kapitel samlar jag ihop och sedan bearbetar, sammanfattar, beskriver och analyserar data för USA rörande hur hushåll allokerat sin tid och sina utgifter. Detta görs av två anledningar. För det första möjliggör det en samlad analys av dessa dimensioner. I det som är andra kapitlet i denna avhandling utvecklar jag en teori, som fokuserar just på hur hushåll fördelar sin tid och sina utgifter, och i den tillägnar jag hemproduktion en viktig roll. Således tjänar detta kapitel, det vill säga kapitel 1, också ett syfte genom att sammanställa ett dataset som kan användas i kapitel 2.

Det är välkänt att hushållen har spenderat en allt större andel av sina totala konsumtionsutgifter på tjänster. Jag finner att denna ökning stämmer överens med utvecklingen bland de olika typer av hushåll som jag undersöker. Vidare så ökade konsumtionen av tjänster i förhållande till icke varaktiga varor under flera årtionden trots ett ökat relativpris på tjänster. Detta är något som vanliga ekonomiska modeller med stabila och homotetiska preferenser inte kan förklara.

I genomsnitt ökade antalet timmar som arbetades i marknaden sedan början av 1960-talet. Detta genomsnitt är dock ett resultat av utvecklingar som ibland skiljde sig markant åt mellan olika grupper.

Det genomsnittliga antalet arbetade timmar minskade bland män, medan det ökade markant bland kvinnor. Medan nedgången bland männen var utbredd, var ökningen bland kvinnor i stor utsträckning koncentrerad till kvinnor i parhushåll. Samtidigt ökade antalet arbetade timmar i hemproduktion bland män, medan de minskade bland kvinnor, speciellt bland kvinnor i parhushåll.

Jag jämför hur timlöner utvecklades för olika typer av män och kvinnor. I genomsnitt steg kvinnors timlöner mer än männens. Genom att granska data närmare, och genom att analysera timlöner tillsammans med arbetade timmar i marknaden respektive i hemmet, uppenbaras vissa mönster som indikerar att utvecklingarna i löner, på egen hand, sannolikt inte kan förklara varför kvinnors arbetsutbud ökat medan det minskat bland män, samt vissa andra mönster kopplade till konsumtion. Exempelvis var timlönerna för kvinnor med den lägre utbildningsnivån i singelhushåll konsekvent lägre än lönenivån bland män med den lägre utbildningsnivån i singelhushåll. Trots detta var utgiftsandelen på tjänster slående lika bland de två grupperna. Vidare var kvinnornas genomsnittliga timlöner konstanta i förhållande till männens under hela perioden som undersöks, men kvinnornas arbetade timmar i marknaden ökade, medan de minskade bland männen. I mitt nästa kapitel lyfter jag fram hemproduktion och sociala normer som potentiellt viktiga faktorer för att kunna förstå skillnader mellan hushåll och utvecklingar över tid kopplade till konsumtion och hur hushåll fördelat sin tid.

Hemproduktion, konsumtionsutgifter och allokeringen av tid

I kapitel 1 visar jag att en stor andel av hushållens tid läggs på hemproduktion. Jag dokumenterar också att skillnaderna mellan hur mycket som kvinnor och män jobbar i hemproduktion har skiljt sig tydligt åt. Om hushållen i viss usträckning kan substituera mellan det som kan köpas i marknaden och det som kan produceras i hemmet, är det möjligt att antalet arbetade timmar i marknaden, och andelen av utgifterna som läggs på olika typer av varor och tjänster, påverkas när relativpriset mellan det som kan köpas i marknaden och det som produceras i hemmet förändras. I detta kapitel undersöker jag i vilken mån det går att förklara de mönster som finns i data med en modell som tar hänsyn till hemproduktion.

Som redan nämndes i sammanfattningen av kapitel 1, finns det vissa mönster i data som en vanlig ekonomisk modell med homotetiska preferenser inte kan förklara. Forskningen har dock kommit långt på att använda icke homotetiska preferenser. I korta drag så innebär dessa att det existerar en inkomsteffekt som gör att hushåll fördelar en större andel av sina utgifter till tjänster när inkomstnivån stiger. Syftet med detta kapitel är att undersöka om en modell som tar hänsyn till hemproduktion men begränsas till användanden av homotetiska preferenser kan förklara tendenserna i data, och i sådant fall vad som krävs.

I modellen konsumerar hushållen tjänster, vilka antigen kan köpas i marknaden eller produceras i hemmet genom att kombinera tid med kapital och icke varaktiga varor. Som ett exempel, behövs det för att producera köttbullar att hushållen kombinerar sin tid med ingredienser (t.ex. köttfärs och kryddor) och kapital (t.ex. spis och stekpanna). För att finansera dessa utgifter på tjänster, icke varaktiga varor och kapital, behöver hushållet ha en inkomst, vilket det får genom att arbeta i marknaden.

Jag fokuserar på kvinnor och män i singelhushåll. En på förhand rimlig hypotes till varför konsumtionen av tjänster har vuxit i förhållande till konsumtionen av icke varaktiga varor skulle kunna vara att hushållen substituerade bort från hemproducerade tjänster, till att köpa dem i marknaden. Jag finner att detta inte stämmer överens med vad som skedde bland männen, som substituerade mot att producera mer tjänster i hemmet. Utvecklingen förklaras av att det blev relativt sett billigare att producera i hemmet, vilket i sin tur delvis berodde på lägre priser på kapital och icke varaktiga varor, och också på att hushållen blev mer effektiva i användandet av dessa produktionsfaktorer. Anlednigen till varför konsumtionen av tjänster ökade i förhållande till konsumtionen av icke varaktiga varor, var att männen minskade kvantiteten av icke varaktiga varor som användes för att framställa tjänster i hemmet.

Kvinnorna substituerade också bort från icke varaktiga varor i hemproduktionen. Men till skillnad från männen, substituterade de från hemproducerade till marknadsproducerade tjänster. Det som ligger bakom denna utveckling är förändringen i vad jag kallar för sociala normer, vilka i min modell antas påverka kvinnornas preferenser för arbete i marknaden relativt arbete i hemmet. Förändringarna i social normer fick kvinnorna att i en allt större utsträckning förflytta sina timmar från hemarbete till marknadsarbete. Utöver att detta fick direkta och betydande konsekvenser för hur kvinnor fördelade sin tid, innebar förändringen i social normer också att kvinnorna började efterfråga mer tjänster, samt att efterfrågan på icke varaktiga varor, som används för att producera tjänster i hemmet, minskade.

Sysselsättningsbeständighet mellan generationer och allokeringen av talang

Det är ett välkänt faktum att många barn väljer samma yrken som sina föräldrar. Genom att använda oss av svenska data kan jag, tillsammans med mina medförfattare John Kramer och Josef Sigurdsson, t.ex. visa att sannolikheten för att en man vars far är läkare själv blir läkare är tio gånger större än vad den är att en annan person blir det. Denna sannolikhetsbias är i genomsnitt sex bland männen som vi studerar.

Hur kommer detta sig? Listan med uppsatser som har ställt liknande frågor är lång. Sedan länge har forskare sökt besvara frågor rörande i vilken utsträckning som diverse utfall beror på genetik och medfödda förmågar, och i vilken utsträckning som de beror på saker som uppfostran. Den fråga som vi ställer oss är i vilken usträckning som söner tenderar att dras mot sina fäders yrken som en följd av att de besitter liknande färdigheter, och i vilken usträckning som det beror på andra faktorer.

Vi finner att yrkesval har stor betydelse för den tydliga positiva korrelationen mellan fäders och söners inkomstranking som existerar i data. Denna slutsats drar vi efter att tilldelat varje son i respektive yrke den genomsnittliga inkomsten i just det yrket och sedan efter detta beräknat korrelationen igen. Resultat är då att korrelationen är i stort sett oförändrad. Således är det viktigt att förstå yrkesval för att kunna förstå intergenerationell inkomstmobilitetet.

Genom att använda oss av data över inkomster och testresultat från färdighetstester som unga män genomgick inför militärtjänstgöringen, uppskattar vi varje individs produktivitet i samtliga av de yrken som vi täcker. Vi finner att förmågor kan förklara en del av den sannolikhetsbias som vi ser i data, men att den största delen måste förklaras av något annat. Men vad detta andra är, är svårt att säga, och kan innefatta allt möjligt, som t.ex. ojämlikhet tillgång till information om yrket och utbilidning, erfarenhet, nätverk och kontakter, eller familjeföretag. Vi bakar ihop dessa faktorer till något som vi kallar för "rabatter", vilket i princip fångar att dessa faktorer underlätter för sönerna att ta sig an samma yrken som sina fäder, jämfört med andra söner. Vi uppskattar storleken på dessa rabatter med hjälp av en strukturell modell och använder sedan modellen för att simulera kontrafaktiska utfall, där vi tar bort rabatterna. Sannolikheten att sönerna väljer samma yrke som sina fäder minskar kraftigt när rabatterna tas bort. Medan cirka 8,6 procent av sönerna väljer samma yrken som sina fäder i ursprungsläget, faller denna andel till 3,4 procent när rabatterna tas bort. Ett intressant resultat är att effekterna på total produktion i ekonomin är försumbara, trots att det är många som gör andra yrkesval. Effekten på korrelationen mellan fäders och söners inkomster faller från 0,245 till 0,217 och förklaras främst av att det är söner till fäder i den allra lägsta delen av inkomstfördelningen som ser sina inkomster stiga.

Som ett komplement till dessa resultat från den strukturella modellen, skattar vi också ett antal empiriska modeller. I dessa mäter vi hur inkomsterna påverkas av att deras fäder har yrken som minskar över tid. När yrkerna blir mindre vanliga, minskar sannolikheten för att sönerna ska välja samma yrken som sina fäder. Idén är sen att uppskatta vilken effekt som detta har på sönernas inkomster. Våra resultat visar att söner som har fäder i yrken som blir mindre vanliga har högre inkomster än vad som annars skulle vara fallet. Vi finner att denna positiva effekt kommer ifrån att söner som är mindre lämpade för sina fäders yrken väljer andra yrken, där de passar bättre in och får en högre inkomst. Vidare är dessa framförallt söner till fäder i den nedre delen av inkomstfördelningen.

Penningpolitik och likviditetsbegränsningar - evidens från euroområdet

Målet bland de flesta av nutidens centralbanker är prisstabilitet, vilket tros minska konjunturfluktuationer och vara gynnsamt för ekonomisk tillväxt. När inflationen är eller förväntas bli för hög, höjer centralbanken styrräntan, för att på så vis minska efterfrågan i ekonomin och därigenom dämpa inflationen. Men hur mycket påverkar dessa förändringar i styrräntan efterfrågan och produktion? Skiljer sig detta mellan länder och i sådant fall varför? Dessa är frågor som jag och mina medförfattare José-Elías Gallegos, John Kramer och Ricardo Lima ställer i detta kapitel.

För att besvara dessa frågor, drar vi nytta av ett det finns ett stort antal länder som alla påverkas direkt av den styrränta som sätts av en gemensam centralbank, nämligen länderna i euroområdet. Att fokusera på euroområdet är för oss idealt. För det första så innebär det att vi endast behöver identifiera *en* uppsättning penningpolitiska chocker. Detta innebär i sin tur att de skillnader som vi uppmäter mellan länderna när det kommer till hur mycket produktionen påverkas, kan tillskrivas just precis att de reagerar olika, och inte att vi identifierat chockerna mer eller mindre väl i de olika länderna. För det andra är länderna, av diverse historiska skäl, i många olika dimensioner olika varandra. Att detta är fallet är såklart nödvändigt för att vi, i ett andra steg, ska kunna undersöka varför effekterna är större i vissa länder än andra.

Vi använder högfrekvent finansiell data för att att identifiera de penningpolitiska chockerna och estimerar sedan effekterna på produktionen i de olika länderna med moderna ekonometriska metoder. Vi finner att produktionen i hög grad påverkas olika i de olika länderna. Genom att framför allt använda oss av data från *Household Finance and Consumption Survey* (HFCS), vilka samlar in och sammanställer data på hushållsnivå om hushållens finanser och konsumtion, sammanställer vi ett dataset med diverse statistiska mått för respektive land. Dessa jämförs sedan med de effekter av penningpolitiska chocker som vi uppmätt.

Flera uppsatser inom den teoretiska litteraturen har tidigare betonat att hushåll med lite likvida medel möjligtvis är viktiga för hur stor effekt ränteförändringar har på ekonomier. Dessa är hushåll som kan vara extra känsliga för inkomstförändringar och som i hög grad behöver anpassa sin konsumtion, medan hushåll med mer likvida medel har större möjlighet att använda sitt sparande för att undvika stora fluktuationer i konsumtionen. Mot bakgrund av ett antal inflytelserika bidrag till den teoretiska litteraturen, beräknar vi andelen av husållen i respektive land som kan klassas som hand-till-mun-hushåll (HtM), som är hushåll med lite likvida medel i förhållande till sina inkomster. Bland alla de statistiska mått som vi undersöker, är det bara HtM som samvarierar med hur mycket produktionen påverkas i länderna. Vi finner att styrränteförändringar påverkar produktionen mer i länder där andelen HtM är högre.

Våra resultat torde vara relevanta för den nationalekonomiska forskningen. Tidigare teoretiska uppsatser har visat att en högre andel HtM-hushåll kan innebära att effekterna av styrränteförändringar blir större, men att det omvända också är möjligt. Våra result kan vägleda den framtida forskningen inom detta område. Våra resultat är också relevanta för beslutsfattare. Specifikt så tyder våra resultat på att centralbanker behöver ta hänsyn till hur stor andel av hushållen som är likviditetsbegränsade när de uppskattar vad som är en väl avvägd penningpolitik. Vidare behöver rådgivare och beslutsfattare på den Europeiska Centralbanken vara varse att de styrränteförändringar som görs har effekter som skiljer sig mellan ekonomierna inom euroområdet.

IIES MONOGRAPH SERIES

1. Michaely, Michael, *The Theory of Commercial Policy: Trade and Protection*, 1973

2. Söderström, Hans Tson, *Studies in the Microdynamics of Production and Productivity Change*, 1974

3. Hamilton, Carl B., *Project Analysis in the Rural Sector with Special Reference to the Evaluation of Labour Cost*, 1974

4. Nyberg, Lars and Staffan Viotti, *A Control Systems Approach to Macroeconomic Theory and Policy in the Open Economy*, 1975

5. Myhrman, Johan, *Monetary Policy in Open Economies*, 1975

6. Krauss, Melvyn, *International Trade and Economic Welfare,* 1975

7. Wihlborg, Clas, *Capital Market Integration and Monetary Policy under Different Exchange Rate Regimes,* 1976

8. Svensson, Lars E.O., *On Competitive Markets and Intertemporal Resources Allocation*, 1976

9. Yeats, Alexander J., *Trade Barriers Facing Developing Countries*, 1978

10. Calmfors, Lars, *Prices, Wages and Employment in the Open Economy*, 1978

11. Kornai, Jànos, *Economics of Shortage, Vols I and II*, 1979

12. Flam, Harry, *Growth, Allocation and Trade in Sweden. An Empirical Application of the Heckscher-Ohlin Theory*, 1981

13. Persson, Torsten, *Studies of Alternative Exchange Rate Systems. An Intertemporal General Equilibrium Approach*, 1982

14. Erzan, Refik, *Turkey's Comparative Advantage, Production and Trade Patterns in Manufactures. An Application of the Factor Proportions Hypothesis with Some Qualifications*, 1983

15. Horn af Rantzien, Henrik, *Imperfect Competition in Models of Wage Formation and International Trade*, 1983

16. Nandakumar, Parameswar, *Macroeconomic Effects of Supply Side Policies and Disturbances in Open Economies*, 1985

17. Sellin, Peter, *Asset Pricing and Portfolio Choice with International Investment Barriers*, 1990

18. Werner, Ingrid, *International Capital Markets: Controls, Taxes and Resources Allocation*, 1990

19. Svedberg, Peter, *Poverty and Undernutrition in Sub-Saharan Africa: Theory, Evidence, Policy*, 1991

20. Nordström, Håkan, *Studies in Trade Policy and Economic Growth*, 1992

21. Hassler, John, Petter Lundvik, Torsten Persson, and Paul Söderlind, *The Swedish Business Cycle: Stylized facts over 130 years,* 1992

22. Lundvik, Petter, *Business Cycles and Growth*, 1992

23. Söderlind, Paul, *Essays in Exchange Rates, Business Cycles and Growth*, 1993

24. Hassler, John A.A., *Effects of Variations in Risk on Demand and Measures of Business Cycle Comovements*, 1994

25. Daltung, Sonja, *Risk, Efficiency, and Regulation of Banks*, 1994

26. Lindberg, Hans, *Exchange Rates: Target Zones, Interventions and Regime Collapses*, 1994

27. Stennek, Johan, *Essays on Information-Processing and Competition*, 1994

28. Jonsson, Gunnar, *Institutions and Incentives in Monetary and Fiscal Policy*, 1995

29. Dahlquist, Magnus, *Essays on the Term Structure of Interest Rates and Monetary Policy*, 1995

30. Svensson, Jakob, *Political Economy and Macroeconomics: On Foreign Aid and Development,* 1996
31. Blix, Mårten, *Rational Expectations and Regime Shifts in Macroeconometrics*, 1997
32. Lagerlöf, Nils-Petter, *Intergenerational Transfers and Altruism*, 1997
33. Klein, Paul, *Papers on the Macroeconomics of Fiscal Policy*, 1997
34. Jonsson, Magnus, *Studies in Business Cycles*, 1997
35. Persson, Lars, *Asset Ownership in Imperfectly Competitive Markets*, 1998
36. Persson, Joakim, *Essays on Economic Growth*, 1998
37. Domeij, David, *Essays on Optimal Taxation and Indeterminacy*, 1998
38. Flodén, Martin, *Essays on Dynamic Macroeconomics*, 1999
39. Tangerås, Thomas, *Essays in Economics and Politics: Regulation, Elections and International Conflict*, 2000
40. Lidbom, Per Pettersson, *Elections, Party Politics and Economic Policy*, 2000
41. Vestin, David, *Essays on Monetary Policy*, 2001
42. Olofsgård, Anders, *Essays on Interregional and International Political Economics*, 2001
43. Johansson, Åsa, *Essays on Macroeconomic Fluctuations and Nominal Wage Rigidity*, 2002
44. Groth, Charlotta, *Topics on Monetary Policy*, 2002
45. Gancia, Gino A., *Essays on Growth, Trade and Inequality*, 2003
46. Harstad, Bård, *Organizing Cooperation: Bargaining, Voting and Control*, 2003
47. Kohlscheen, Emanuel, *Essays on Debts and Constitutions*, 2004
48. Olovsson, Conny, *Essays on Dynamic Macroeconomics*, 2004

49. Stavlöt, Ulrika, *Essays on Culture and Trade*, 2005
50. Herzing, Mathias, *Essays on Uncertainty and Escape in Trade Agreements*, 2005
51. Bonfiglioli, Alessandra, *Essays on Financial Markets and Macroeconomics*, 2005
52. Pienaar, Natalie, *Economic Applications of Product Quality Regulation in WTO Trade Agreements*, 2005
53. Song, Zheng, *Essays on Dynamic Political Economy*, 2005
54. Eisensee, Thomas, *Essays on Public Finance: Retirement Behavior and Disaster Relief*, 2005
55. Favara, Giovanni, *Credit and Finance in the Macroeconomy*, 2006
56. Björkman, Martina, *Essays on Empirical Development Economics: Education, Health and Gender*, 2006
57. Larsson, Anna, *Real Effects of Monetary Regimes*, 2007
58. Prado, Jr., Jose Mauricio, *Essays on Public Macroeconomic Policy*, 2007
59. Tonin, Mirco, *Essays on Labor Market Structures and Policies*, 2007
60. Queijo von Heideken, Virginia, *Essays on Monetary Policy and Asset Markets*, 2007
61. Finocchiaro, Daria, *Essays on Macroeconomics*, 2007
62. Waisman, Gisela, *Essays on Discrimination and Corruption*, 2008
63. Holte, Martin Bech, *Essays on Incentives and Leadership*, 2008
64. Damsgaard, Erika Färnstrand, *Essays on Technological Choice and Spillovers*, 2008
65. Fredriksson, Anders, *Bureaucracy, Informality and Taxation: Essays in Development Economics and Public Finance*, 2009
66. Folke, Olle, *Parties, Power and Patronage*, 2010

67. Drott, David Yanagizawa, *Information, Markets and Conflict*, 2010
68. Meyersson, Erik, *Religion, Politics, and Development*, 2010
69. Klingelhofer, Jan, *Models of Electoral Competition*, 2010
70. Perrotta, Maria Carmela, *Aid, Education and Development*, 2010
71. Caldara, Dario, *Essays on Empirical Macroeconomics*, 2011
72. Mueller, Andreas, *Business Cycles, Unemployment and Job Search*, 2011
73. von Below, David, *Essays in Climate and Labour Economics*, 2011
74. Gars, Johan, *Essays on the Microeconomics of Climate Change*, 2012
75. Spiro, Daniel, *Some Aspects of Resource and Behavioral Economics*, 2012
76. Ge, Jinfeng, *Essays on Macroeconomics and Political Economy*, 2012
77. Li, Yinan, *Institutions, Political Cycles and Corruption*, 2012
78. Håkanson, Christina, *Changes in Workplaces and Careers*, 2013
79. Qin, Bei, *Essays on Empirical Development and Political Economics*, 2013
80. Jia, Ruixue, *Essays on the Political Economics of China's Development*, 2013
81. Campa, Pamela, *Media Influence on Pollution and Gender Equality*, 2013
82. Seim, David, *Essays on Public, Political and Labor Economics*, 2013
83. Shifa, Abdulaziz B., *Essays on Growth, Political Economy and Development*, 2013
84. Panetti, Ettore, *Essays on the Economics of Banks and Markets*, 2013

85. Schmitt, Alex, *Beyond Pigou: Climate Change Mitigation, Policy Making and Decisions*, 2014
86. Rogall, Thorsten, *The Economics of Genocide and War*, 2015
87. Baltrunaite, Audinga, *Political Economics of Special Interests and Gender*, 2016
88. Harbo Hansen, Niels-Jakob, *Jobs, Unemployment, and Macroeconomic Transmission*, 2016
89. Stryjan, Miri, *Essays on Development Policy and the Political Economy of Conflict*, 2016
90. Karadja, Mounir, *On the Economics and Politics of Mobility*, 2016
91. Kitamura, Shuhei, *Land, Power and Technology, Essays on Political Economy and Historical Development*, 2016
92. Malmberg, Hannes, *Human Capital in Development Accounting and Other Essays in Economics*, 2017
93. Öberg, Erik, *On Money and Consumption*, 2017
94. Lane, Nathaniel, *States of Development: Essays on the Political Economy of Development in Asia*, 2017
95. Prawitz, Erik, *On the Move: Essays on the Economic and Political Development of Sweden*, 2017
96. Iwanowsky, Mathias, *Essays in Development and Political Economics*, 2018
97. Dehdari, Sirus Håfström, *Radical Right, Identity, and Retaliation*, 2018
98. Harmenberg, Karl, *Essays on Income Risk and Inequality*, 2018
99. Kilström, Matilda, *Households' Responses to Policy in Labor and Credit Markets*, 2018
100. Meriläinen, Jaakko, *Essays in Political Economics*, 2019
101. Sigurdsson, Jósef , *Essays on Labor Supply and Adjustment Frictions*, 2019
102. Cocciolo, Serena, *Participatory Governance and Public Service Provision* , 2019

103. Mitrunen, Matti, *Essays on the Political Economy of Development*, 2019

104. Olsson, Jonna, *Work, wealth, and well-being: Essays in macroeconomics*, 2019

105. Darougheh, Saman, *Search and Mismatch*, 2019

106. van Vlokhoven, Has, *On the Cost of Capital, Profits and the Diffusion of Ideas*, 2020

107. Foltyn, Richard, *Essays in Macroeconomics and Household Finance*, 2020

108. Lerva, Benedetta, *Worms, Farms and Schools: Three Essays on Investment Choices in Uganda*, 2020

109. Kragh-Sørensen, Kasper, *On Housing, Mortgages, and Taxation*, 2020

110. Karlman, Markus, *Essays on Housing: tax treatment, prices, and macroeconomic implications*, 2020

111. Kinnerud, Karin, *Financial Choice and Public Policy*, 2020

112. Ghisolfi, Selene, *Fairness, technology adoption, water sanitation and pandemic control: Six essays on four topics in Development Economics*, 2020

113. Åhl, Magnus, *Essays on Individual-Level Wage Stickiness and Forward Guidance*, 2020

114. Dev, Divya *Essays on Gender, Development and Political Economy*, 2021

115. Gallegos Dago, José Elias *Essays on Macroeconomics*, 2022

116. Sun, Xueping *Essays on China's Economic Development: Innovation, Public Debt and Social Connections*, 2022

117. Kramer, John Vincent *Essays on Macroeconomics, Monetary Policy and Mobility*, 2022

This dissertation contains four chapters. The first two chapters analyze and try to understand how households in the U.S. have spent their time, and how they have allocated their expenditures on different types of consumption. In the third chapter, the importance of family background for occupational choice and its implications for intergenerational earnings mobility is studied. The fourth chapter estimates differential effects of monetary policy on output in countries in the euro area and investigates potential explanations for it.

Mattias Almgren
holds an M.Sc. in Economics
from Uppsala University

ISBN 978-91-8014-270-0
ISSN 0346-6892

Department of Economics

Stockholm University