Essays on Dynamics of Household and Firm Choices

by

Aruni Mitra

B.Sc., University of Calcutta, 2012
M.S., Indian Statistical Institute, 2014

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate and Postdoctoral Studies

(Economics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

July 2020

The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the dissertation titled:

Essays on Dynamics of Household and Firm Choices

submitted by **Aruni Mitra**

in partial fulfillment of the requirements for the degree of **Doctor of Philosophy** in **Economics**.

Examining Committee:

Giovanni Gallipoli, Professor, Economics, UBC
Supervisor

Michael B. Devereux, Professor, Economics, UBC
Supervisory Committee Member

Henry E. Siu, Professor, Economics, UBC
University Examiner

Lorenzo Garlappi, Professor, Business Administration, UBC
University Examiner

Thomas F. Crossley, European University Institute
External Examiner

Additional Supervisory Committee Member:

Florian Hoffmann, Associate Professor, Economics, UBC
Supervisory Committee Member

Abstract

For nearly four decades in the post-War United States, productivity rose during economic booms and fell in recessions. The *first chapter* of this thesis studies how increased labour market flexibility because of rapid de-unionization since the early 1980s can explain the sudden vanishing of this procyclicality of productivity, the so-called 'productivity puzzle'. Falling costs of hiring and firing workers, due to the decline in union power, prompted firms to rely more on employment adjustment (extensive margin) instead of changing workers' effort through labour hoarding (intensive margin). High dependence on labour hoarding explains productivity's historical procyclicality, and its reduced importance in recent decades explains why productivity is now less procyclical. Increased hiring and firing of workers also imply a rise in the relative volatility of employment. I show that U.S. states and industries with a larger drop in union density experienced a deeper fall in the procyclicality of productivity and a larger increase in the relative volatility of employment.

Simultaneous to the productivity puzzle in the mid-1980s, there were other important structural changes in the U.S. economy, namely, the rise of the service sector, increased use of intangible capital, more accommodative monetary policy, and the decline in the volatility of shocks during Great Moderation. The *second chapter* shows that none of these structural changes can explain the productivity puzzle. However, allowing the hiring cost to decline between pre- and post-1980s in an otherwise standard New Keynesian model with endogenous effort can match almost all the fall in cyclical productivity correlations and the rise in the relative volatility of employment.

The *third chapter* characterizes the joint evolution of cross-sectional inequality in permanent income and consumption among parents and children in the U.S. We use a model of intra-family persistence across generations to estimate the parameters determining inequality of consumption and income within a generation. In accounting for cross-sectional dispersion, we find that idiosyncratic heterogeneity is quantitatively more important than inequality arising from family factors. This suggests that parents provide limited insurance against idiosyncratic life-cycle risk, even though the levels of permanent income and consumption exhibit significant persistence across generations.

Lay Summary

My dissertation consists of three chapters studying the business cycle dynamics of productivity and the evolution of intergenerational inequality in the United States. In the first chapter, I argue that easier hiring and firing of workers due to the rapid decline in labour-union power can explain why productivity suddenly started to rise during recessions since the mid-1980s, after having risen during economic booms for almost four decades. In the second chapter, I rule out structural changes that occurred in the U.S. economy around the same time, like the rise of the service sector, the increased use of intangible capital, or the change in monetary policy, as significant explanations for driving changes in productivity dynamics. The third chapter studies how inequality in income and consumption evolves over generations and finds that an overwhelming majority of the observed inequality cannot be explained by intra-family linkages.

Preface

Chapters 2 and 3 of the thesis are pieces of original, unpublished and independent work. Chapter 4 titled "Consumption and Income Inequality across Generations" is ongoing collaborative work with Professor Giovanni Gallipoli (Vancouver School of Economics, University of British Columbia) and Professor Hamish Low (Nuffield College, University of Oxford). I have been involved throughout each stage of the research project: preparing data for statistical analysis, designing and executing the estimation strategy, and writing and editing the manuscript.

Table of Contents

List of Tables

List of Figures

Acknowledgements

I would like to express my deepest gratitude to my supervisor, Giovanni Gallipoli for his support and advice throughout my doctoral studies. I am greatly indebted to Michael Devereux for his insightful comments and suggestions which have immensely enriched my learning experience. I also want to thank Paul Beaudry, who, even though not formally on my thesis committee, devoted a lot of time to discuss and shape my research in its early stages. Hamish Low, with whom a chapter of this thesis is co-authored, has been a constant source of encouragement and motivation. I also thank him for inviting me for a research visit to the Institute for New Economic Thinking at the University of Cambridge during my PhD. I am also indebted to Nicole Fortin, David Green, Florian Hoffmann, Thomas Lemieux, Jesse Perla, Henry Siu and other participants of the empirical economics and macroeconomics brown-bag seminars at the Vancouver School of Economics (VSE) for their very useful feedback and discussions.

I have been extremely fortunate to have some wonderful fellow students at VSE. Davide Alonzo, Anujit Chakraborty, Anand Chopra, Matthew Courchene, Tsenguun Enkhbaatar, Nicolas Franz-Pattillo, Arkadev Ghosh, Nadhanael GV, Da Kang, Neil Lloyd, David Macdonald, Bipasha Maity, Benjamin Milner, Ronit Mukherji, Adlai Newson, Mengying Wei and Michael Wiebe — they have all helped me more than they know themselves.

I want to thank Prithu Banerjee, Taylor Chapman, Amartya Dutta, Alexander Firanchuk, Bert Kramer and Jonathan Schmok — friends outside my PhD program, for their companionship that kept me sane at various points of time during my six-year-long journey of doctoral studies. We had some wonderful time together.

I would like to thank my professors at St. Xavier's College and Indian Statistical Institute in Kolkata, who had taught how to see the world around me through the lens of economics. Two of my teachers from high school days, Supriya Datta and Kalisadhan Mukherjee, both of whom passed away last year, deserve special mention. Without their motivation, encouragement and nurturing of my young mind, I would not be where I am today.

Finally, I thank my mother, who had given up her doctoral studies to raise me, and the memory of my father, who would have been happy to see this thesis in print. Without their love, sacrifice and patience, this thesis would not be a reality.

Chapter 1

Introduction

This thesis studies the dynamics of household and firm choices. The following three chapters touch upon areas of business cycle dynamics of productivity as well as long-run evolution of economic inequality across generations. In particular, the first chapter studies the impact of the decline in the power of labour unions in influencing firms' hiring and firing decision, which in turn affects the business cycle movements of productivity. The second chapter looks at a host of structural changes in the U.S. economy around the mid-1980s, and once again studies the impact of these significant regime-switches on the business cycle properties of productivity. Finally, the third chapter moves beyond the short-run view of the economy at business cycle frequencies and instead looks at how economic inequality propagates through intra-family linkages across generations.

Productivity, as measured by either the total value-added per worker or the total output produced per hour worked or even the total factor productivity (TFP), rose during economic booms and fell in recessions for about four decades after World War II. However, around the mid-1980s, this procyclicality of productivity suddenly started to vanish. Productivity has remained countercyclical with employment and total hours worked and acyclical with output ever since. This phenomenon is known as the *productivity puzzle* (see McGrattan and Prescott, 2012).[1] Chapter 2 of this thesis argues that easier hiring and firing of workers due to the rapid decline in union power from the early 1980s in the U.S., precipitated by President Reagan's deregulation measures, can explain this puzzle.

When for some reason the hiring and firing of workers become easier, firms do not need to rely on making their available workers change their effort level depending on economic slack or boom. They can instead simply fire workers in recessions and hire them back during better times. This means that an individual worker's productivity does not need to fluctuate along the business cycle when hiring and firing costs are low. This is a candidate theoretical explanation of the sudden vanishing procyclicality of productivity — the so-called reversal of the *labour hoarding* phenomenon that was cited widely for the historical procyclicality of productivity since the 1940s. Chapter 2 first provides empirical evidence for this channel of a drop in labour hoarding and then identifies de-unionization in the U.S. as a potential reason for the decline in employment adjustment friction.

There are two ways through which the reduced importance of labour hoarding is empirically

[1] In popular media, the term *productivity puzzle* has recently been used in different contexts to mean a variety of phenomena in the U.S. economy, e.g., the slow growth of productivity in recent years, the divergence between labour productivity and real wage growth, etc. However, most of the academic literature recognizes the vanishing procyclicality of productivity as the *'productivity puzzle'*, and it is in this sense that the term will be used in this thesis.

made apparent. First, using a widely accepted decomposition of the TFP measure (see Fernald, 2014) into a measure of factor utilization rate (which arguably measures the intensive margin of factor use) and a utilization-adjusted *true* productivity component, it is shown that the entire drop in procyclicality of TFP is driven by the decline in the explanatory power of the highly procyclical factor utilization component, while the true productivity dynamics changed very little. Secondly, following the strategy in Galí and Gambetti (2009), a structural vector autoregression (SVAR) using productivity and per capita hours shows that it is not the changing relative importance of the technology versus demand shocks but the changing response of the U.S. economy to the same shocks before and after the 1980s that led to the fall in productivity correlations with output and labour input. In particular, it is shown that conditional on the same expansionary demand shock in the pre- and post-1984 eras, productivity rose much mildly in the post-1984 period. This can be seen as indirect evidence of firms relying more on hiring new workers to meet the extra demand instead of increasing the productivity (or effort) of its available workers due to the drop in hiring cost. Moreover, this enhanced hiring and firing also translate into larger volatility of employment relative to both output and the intensive margin of factor use at the same time as the productivity puzzle.

After considering a variety of reasons as to why hiring and firing costs can decline, e.g., the advent of online job search, increased use of part-time or temporary workers, I settled for de-unionization as the most plausible channel.[2] The de-unionization episode in the U.S. not only lines up well in terms of the timing and speed of the productivity puzzle but it also aligns well with international evidence, e.g., Canada experienced neither de-unionization nor the productivity puzzle, while the United Kingdom under Margaret Thatcher experienced both a decline in union clout and the vanishing procyclicality of productivity. Using cross-sectional variation from 51 U.S. states and 17 industries, it is argued that sectors and regions which experienced a larger decline in the fraction of unionized workers also witnessed a larger drop in the procyclicality of productivity and a bigger increase in the volatility of employment relative to that of output. Moreover, it is also shown how employment protection laws like the *right-to-work* laws in U.S. states interacts with the de-unionization. In particular, those states which had these right-to-work laws understandably did not get impacted by the de-unionization episode (as they were already pro-business) and hence did not observe the productivity puzzle, while the entire effect of the decline of union power was concentrated in the non-right-to-work states.[3]

The idea of a greater labour market flexibility in the last three decades in the U.S. as an explanation for the vanishing procyclicality of productivity was first discussed in Gordon (2011), and most recently in Galí and van Rens (2017). The key insight of the latter is that a decline in labour market turnover, which reduced hiring frictions, can match the observed decline in labour productivity correlations. However, Galí and van Rens (2017) does not pinpoint any particular

[2]See Stansbury and Summers (2020) for a survey of the changes in the U.S. economy that were brought about by the rapid loss of workers' bargaining power from the early 1980s.

[3]Autor (2003) shows how in heavily unionized states hiring had become so difficult that it had to be outsourced to temporary workers.

structural reason as to why the labour market turnover suddenly changed in the mid-1980s. Chapter 2 investigates various structural changes in the labour market as possible candidates for the falling hiring and firing frictions, and finds that de-unionization is the most likely channel.[4]

Chapter 3 extends the understanding of the productivity puzzle by considering potential alternative explanations for the puzzle and quantifying the importance of the hiring cost channel identified in Chapter 2 as a feasible explanation.

McGrattan and Prescott (2012) propose that under-estimating the use of intangible capital is the main channel for explaining the vanishing procyclicality of productivity. They argue that if intangible capital is strongly procyclical but it is not included while measuring output, then the measured procyclicality of productivity with value-added will be less. However, they do not present any empirical fact regarding the increased use of intangible capital around the mid-1980s, rather focusing on the Great Recession period of the late 2000s. I do not find cross-industry evidence that greater investment in intellectual property products (which is the closest measurable proxy for intangible capital) is correlated with a larger drop in productivity correlations in the mid-1980s.

Moving beyond the narrative of mismeasurement in factor input (see Galí and van Rens, 2017) or output (see McGrattan and Prescott, 2012), Barnichon (2010) shows that a large portion of the rise in the correlation of labour productivity with unemployment is accounted for by the increasing volatility of technology shocks relative to demand shocks. The main argument is that since technology shocks induce countercyclicality of productivity with labour input and demand shocks induce procyclicality of productivity with labour input, the increasing relative importance of technology shocks will help to explain the fall in unconditional correlations of productivity with hours. However, this explanation is insufficient for two reasons. First, the productivity correlations fell even conditional on a demand shock, whereas conditional on technological shock, the correlation of labour productivity increased slightly in the post-1984 period (as shown in Chapter 2). These changes in the conditional correlations indicate that some form of structural change is required to fully explain the productivity puzzle. Second, the correlation of labour productivity with output conditional on a technology shock has always been positive, and it increased further after the mid-1980s. Hence, the increasing importance of technology shocks cannot explain the falling correlation of output with productivity. Van Zandweghe (2010) performs a comparative study of the changes in technology and demand shocks on one hand and the structural change in the labour market on the other. He concludes that since the productivity correlations conditional on both demand and supply shocks have changed, it is more likely that change in labour market flexibility is the key factor behind the phenomenon. Through the lens of a New Keynesian model, featuring endogenous effort choice and costly hiring of workers, I show in Chapter 3, that neither the reduced volatility during Great Moderation nor the changing relative importance of technology and demand shocks

[4] The model in Galí and van Rens (2017) predicts that productivity co-moves positively with both output and inputs in response to technology shocks. This is in stark contrast to the empirical finding in the post-War U.S. where labour input is negatively correlated with technology shocks (see for example Galí and Gambetti (2009)). A key element of the model in Chapter 3 will be to generate impulse responses to technology and demand shocks that mimic the empirically observed ones.

can explain any notable part of the puzzle while allowing the hiring cost to fall by 67% to reflect the decline in private-sector union density can generate almost the entire drop in productivity correlations and more than 80% of the increase in the relative volatility of employment.

In contrast to studying demand and supply shocks, Garin, Pries, and Sims (2018) look at changes in the contribution of sector-specific shocks relative to aggregate shock in the total volatility of industrial production, and claim that the increasing importance of sectoral shocks have caused a drop in productivity correlations and also led to jobless recoveries.[5] However, I show that the empirical finding of the increasing importance of sectoral shocks is not robust to the choice of the dataset, and there is reason to doubt that there has been a sudden increase in sectoral reallocation towards more productive sectors since the mid-1980s. Relating jobless recoveries to the vanishing procyclicality of productivity is also done by Berger (2016), whose model relies on heterogeneity in worker productivity. Crucially, in his model, firms know individual worker productivity and during streamlining and restructuring in a recession they can lay off their least productive workers, thereby bringing up the average labour productivity. I show in Chapter 3 that while it is true that firms indeed fire their least productive workers during recessions, this phenomenon has not increased in intensity after the mid-1980s. So, rather than relying on selectively firing the low productivity workers, I find that firms have been hiring and firing all workers more frequently. Moreover, since the focus in Berger (2016) is only on the recovery part of the business cycle, that model has trouble in generating the full magnitude of the drop in productivity correlations. Also, the impulse responses of employment to a TFP shock is positive in Berger's modified Real Business Cycles model, which goes against the empirical evidence and is corrected in the New Keynesian framework I consider.

The link between the diminishing procyclicality of labour productivity and jobless recoveries cannot be overlooked. Even though jobless recoveries correspond to only a fraction of the entire business cycle (namely, the post-recession recovery and not the pre-recession boom or the recession itself), they are consistent with the falling correlation between productivity and labour input. However, jobless recoveries are at odds with the falling correlation between output and productivity. Hence, there is no a priori reason to try to explain these two phenomena of jobless recoveries and productivity puzzle jointly. Nevertheless, it is natural to ask whether factors explaining one can also explain other. Panovska (2017) runs a horse race among the different channels of explanations for jobless recoveries — (i) reallocation of resources across sectors (like in Groshen and Potter (2003), Garin, Pries, and Sims (2018) and others) or occupations (as highlighted by Jaimovich and Siu (2015)), (ii) expansionary overhang which leads to a restructuring in recessions, like in Berger (2016), (iii) shorter duration of recessions, as pointed out by Bachmann (2012), and (iv) structural changes in the economy, like those highlighted by Galí and Gambetti (2009). Using a VAR analysis, she finds that structural change in the economy is the most plausible channel explaining jobless recoveries. Therefore, in this thesis, I will focus exclusively on the structural

[5]The last three recessions in the U.S. (since the early 1990s recession) were characterized by recoveries where output and productivity picked up (albeit at a slower pace than previous recessions) but labour input did not rise in the initial recovery phase. This phenomenon has been termed as *jobless recoveries*.

changes in the economy to explain the productivity puzzle. Finally, Brault and Khan (2020) point out changes in non-contemporaneous correlations of productivity with output and labour input, and argue that none of the theoretical explanations explored in the literature for the productivity puzzle can explain these changes in non-contemporaneous correlations. Investigating this aspect of cyclical productivity dynamics remains an agenda for future research.

Unlike the first two chapters which deal with short-run business cycle dynamics of firms' choices regarding employment and labour effort, Chapter 4 addresses an issue of long-run dynamics in the economy, namely, the evolution of inequality across generations of households. It asks how much of the observed cross-sectional inequality in earnings, other income sources (like transfers, wealth income and spousal income) and consumption expenditure can be explained by intra-family intergenerational linkages.

Parents influence their children's life-cycle outcomes in many ways. Economists often quantify these influences using measures of intergenerational persistence along dimensions of heterogeneity such as earnings, wealth, or consumption.[6] The various channels of family influence are inter-related as parents can affect their children's outcomes in complex ways: through choices about education, through the transmission of ability and preferences, by providing income-enhancing opportunities, as well as through inter-vivos and bequest transfers affecting wealth and consumption.[7] Further, these mechanisms may be substitutes: investing in a child's education to increase their earnings potential may imply lower transfers of wealth. This rich set of influences suggests that changes in a household's financial circumstances may induce intergenerational effects with empirically testable implications. Unlike most previous studies, which looked at either income or consumption in isolation, and focused primarily on the intergenerational pass-through parameters, we develop a parsimonious model of the *joint* persistence of expenditures, earnings and other income and focus directly on understanding inequality. The co-dependence of these processes turns out to be of critical importance, as discussed in Alan, Browning, and Ejrnæs (2018). Our work has two main objectives: first, to examine the diverse ways parental influences shape children's economic outcomes in a unified framework; second, to quantify how much of the inequality observed in a particular generation is due to parental factors.

To describe the influence of parental heterogeneity, we model intergenerationally linked households that make consumption and saving choices in an environment where persistent shocks shape permanent income. In our baseline model, we characterize the distribution of endogenous expendi-

[6] Research linking family outcomes across generations often focuses on income and earnings persistence (for a survey, see Aaronson and Mazumder, 2008). Related work documents the persistence of wealth (e.g., Charles and Hurst, 2003), consumption (e.g. Waldkirch, Ng, and Cox, 2004; Charles, Danzinger, Li, and Schoeni, 2014; Bruze, 2016) and occupations (Corak and Piraino, 2010; Bello and Morchio, 2016). Boar (2017) documents parental behaviours consistent with a precautionary motive geared to insure children against life-cycle risk.

[7] For the role of transfers, see Daruich and Kozlowski (2016), Bolt, French, Maccuish, and ODea (2018) and Abbott, Gallipoli, Meghir, and Violante (2019). Restuccia and Urrutia (2004) and Lee and Seshadri (2019) examine the role of investments at different stages of the life-cycle. Caucutt and Lochner (2019) highlight the impact of credit constraints in an environment with sequential parental investments. Gayle, Golan, and Soytas (2018) find that human capital accumulation, nonlinear returns to hours worked, and parental time investments are key for the intergenerational correlation in earnings through their effects on fertility and the division of labour.

tures alongside a standard income process. The intergenerational linkages stem from intra-family persistence of earned income as well as from savings and transfer decisions. Specifically, we allow parents to influence outcomes of children in three ways: through earned income, through other sources of income such as transfers, and through consumption.

A key contribution of our analysis is to document the role of parental factors in explaining inequality in the next generation and to contrast the importance of family heterogeneity with the impact of idiosyncratic variation which is independent of parents. The extent to which inequality among parents is passed through to inequality among children depends on intergenerational elasticities; however, the relative importance of family factors for inequality among children also depends on the magnitude of idiosyncratic (family-independent) variation. Hence, a decomposition of observed inequality requires estimates of intergenerational pass-through parameters, estimates of inequality among parents and estimates of idiosyncratic heterogeneity. We use our model to jointly estimate, through a method-of-moments approach, the parameters determining the importance of the different components of inequality.

The model delivers moment restrictions on the variances of earnings, other income and consumption of parents and their adult children, and on their covariation across generations. We use these moments to jointly estimate the parameters dictating intergenerational linkages, as well as the responses of income and consumption to different shocks. Then, through the model, we quantify the contribution of parental factors to children's outcomes and overall cross-sectional inequality.

For estimation, we employ data from the Panel Study of Income Dynamics (PSID) covering birth-cohorts of individuals born between the early 1950s and the late 1970s. We link households' income, expenditures and other family characteristics across generations in a long panel format.[8] To avoid the selection issues associated with women's labour force participation, we focus on a sample of father-son pairs to characterize earnings persistence; however, we include women's labour earnings within our measure of other income. Since the availability of expenditure data varies across survey waves, research on the consumption pass-through based on PSID data has used food expenditures for the full-length sample or restricted attention to shorter periods for which extensive consumption records are available.[9] To account for these data limitations, in our baseline estimation we use information about the higher moments of measured food consumption going back to the late 1960s; then, in a set of robustness checks, we document the robustness of our findings by replicating the analysis with imputed measures of total household outlays,[10] and by restricting the sample period to have detailed expenditure records for most categories.

[8]The PSID initially recorded only housing and food-related expenditures. After 1999 more consumption categories were added to the survey; since 2005, the PSID expenditure data cover all the categories in the Consumption Expenditure Survey (CEX). The CEX started providing detailed data about multiple consumption categories at the micro-level in the 1980s, yet that data is unsuitable for intergenerational studies because individuals were followed for a maximum of four quarters.

[9]The former approach is that of Waldkirch, Ng, and Cox (2004). The latter route is taken by Charles, Danzinger, Li, and Schoeni (2014), who project average consumption expenditures of adult children for the years 2005-2009 on similar measures for their parents in the same years.

[10]The imputation method (see Attanasio and Pistaferri, 2014) relies on the estimation of a demand system using the rich set of expenditures reported post-1999 in the PSID.

Our analysis also highlights a negative association between income inequality and economic mobility. This arises because greater income inequality in the kids' generation is explained by stronger intergenerational pass-through channels. The finding is consistent with the empirical observation that more unequal societies exhibit lower earnings mobility across generations, a relationship often dubbed the 'Great Gatsby' curve (Krueger, 2012; Corak, 2013). The presence of a negative association between mobility and inequality does not, however, pin down the direction of causality (see Rauh, 2017; Comerford, Rodriguez Mora, and Watts, 2017; Kanbur, 2019). While higher persistence can lead to higher inequality (a channel which we explore directly), higher inequality itself might skew the distribution of economic opportunities and stifle long-term mobility. The negative correlation between inequality and intergenerational mobility also does not imply that a decline in mobility is necessary for the rise in inequality. We show that while inequality has increased in the U.S. over the past few decades, mobility has remained roughly stationary, thereby implying a rise in the idiosyncratic life-cycle shocks to the younger generation that cannot be attributed to intra-family linkages.

Our semi-structural approach establishes that idiosyncratic life-cycle heterogeneity, rather than parental economic circumstances, accounts for most of the observed cross-sectional inequality. However, our statistical model remains silent about the underlying sources of these idiosyncratic heterogeneities and intergenerational persistence. Endogenizing intergenerational linkages in a heterogeneous agent framework to bring forth these underlying sources remains as work for future research.

Chapter 2

The Productivity Puzzle and the Decline of Unions

2.1 Introduction

For almost half a century after World War II, labour productivity was procyclical in the United States — it rose during economic booms and fell in recessions. However, since the mid-1980s, it became acyclical with output and countercyclical with total hours worked.[11] Quite strikingly, during the recent Great Recession of the late 2000s, while output and hours took a downward turn, labour productivity stayed constant or even increased slightly over some quarters (Mulligan, 2011). This change in the cyclical correlations has been well documented and is often referred to as the '*labour productivity puzzle*' (McGrattan and Prescott, 2012).

Typically, the explanation for procyclical labour productivity has been the phenomenon of '*labour hoarding*'[12], whereby firms use their available workers less intensely during economic downturns, and more intensely during booms. Since such changes in the intensity of factor utilization cannot be observed in the changes of actual employment or labour hours, the measured labour productivity, which is defined as output per hour worked, appears to be procyclical.[13] Therefore, when the procyclicality of labour productivity started to diminish in the mid-1980s, a natural candidate for an explanation was the vanishing procyclicality of factor utilization. Now, firms resorted to labour hoarding because it was costly for them to hire and fire workers along the business cycle. Therefore, any explanation for reduced labour hoarding should involve more flexible labour market institutions which brings down the cost of employment adjustment. Such a reduction in labour market frictions can also explain the steady increase in the volatility of employment relative to that of output since the mid-1980s. This chapter identifies rapid de-unionization since the early 1980s as a key factor behind increased labour market flexibility in the U.S. that made hiring and

[11] The level of correlation depends on the choice of the filtering process to extract the cyclical part from raw time-series data, but the fall in productivity correlations is robust.

[12] Biddle (2014) notes that the concept of 'labour hoarding', at least in its modern form, dates back to Okun (1962). Burnside, Eichenbaum, and Rebelo (1993) finds that a significant proportion of movements in standard productivity measures like the Solow residual are artefacts of labour hoarding behaviour. By the 1980s, the concept was being regularly used as a standard textbook explanation for procyclical labour productivity (e.g., Dornbusch and Fischer, 1981; Hamermesh and Rees, 1984). Paradoxically, it was around the same time that labour productivity started being more countercyclical.

[13] Real business cycle (RBC) models differ on the explanation of productivity procyclicality. They argue that business cycles are driven by procyclical technology shocks. In Section 2.3.2, I show evidence of negative response of labour inputs to positive technology shocks, which militates against the RBC paradigm.

firing workers easier for firms. In a cross-section of U.S. states and industries, the magnitude of the decline in union density is shown to be significantly correlated with the drop in labour productivity correlations and the rise in the relative volatility of employment. Other structural changes in the labour market like the increased use of part-time workers and the rise in online job search market do not appear to have caused a drop in productivity correlations.

Declining hiring costs should be accompanied by changes in how the economy responds to different types of shocks. For example, in response to a positive demand shock, firms can now increase their labour input by hiring more workers, and hence do not need to increase the intensity of labour utilization by as much. This would imply that the improvement in measured labour productivity and total factor productivity (TFP) in response to a positive demand shock will be significantly reduced. Using a time-varying structural VAR analysis, I show that this is indeed the case.

The rest of the chapter is organized as follows. Section 2.2 documents that the cyclical correlations of productivity with both output and labour input have decreased quite abruptly around the mid-1980s in the U.S. After showing that this puzzling finding is robust to different data sources, the choice of de-trending methodology, and even the measure of labour input, I turn to explain the puzzle in Section 2.3. I investigate the issue in two different ways. First, in Section 2.3.1, I decompose productivity into factor utilization rate and utilization-adjusted productivity and show that it is only the factor utilization component of measured productivity that has become more countercyclical. Second, in Section 2.3.2, I show that the response of the aggregate U.S. economy to technology and demand shocks have changed around the mid-1980s. Both these findings point towards a structural change in the labour market that made hiring and firing workers suddenly much less costly. In Section 2.4, I consider the possible reasons for the drop in employment adjustment cost and show that de-unionization is one structural change in the labour market that is consistent both in terms of timing and speed with the productivity puzzle. Cross-sectional evidence from U.S. states and industries, as well as international evidence from OECD countries, point towards a decline in union power as a major contributing factor towards higher labour market flexibility that explains the productivity puzzle. Finally, Section 2.5 concludes the chapter.

2.2 The Productivity Puzzle

The *productivity puzzle* refers to the sudden vanishing of procyclicality of productivity around the mid-1980s in the U.S. The existing literature on this puzzle has typically used average labour productivity, defined as output per hour worked, as the measure of productivity. In Panels (a) and (b) of Figure 2.1, I corroborate that finding using quarterly data on output and hours worked for the U.S. business sector from 1947 through 2017, sourced from the Labor Productivity and Costs (LPC) dataset of the Bureau of Labor Statistics (BLS). As an alternative measure of productivity, in Panels (c) and (d), I use TFP (unadjusted for factor utilization), sourced from Fernald (2014), and find a remarkably similar pattern of a sudden drop in contemporaneous productivity correlations.[14] While I have used the Baxter and King (1999) bandpass filter to extract the cyclical component of the time-series variables in Figure 2.1, the finding is robust to the choice of the statistical filter.[15]

These changes in productivity correlations have implications for the co-movement of productivity with job flows over the business cycle. One can think of changes in employment as being composed of an inflow of workers through job creation or vacancies, and outflow through job separations. In fact, one can write employment growth Δn_t as the difference between job creation h_t and job destruction f_t, $\Delta n_t = h_t - f_t$, implying $Cov\,(\Delta n_t, lp_t) = Cov\,(h_t, lp_t) - Cov\,(f_t, lp_t)$, where lp_t denotes cyclical component of labour productivity. Then it is natural to expect that job-creation or vacancy rate should become more countercyclical, and job-destruction or separation rate more procyclical after the 1980s. In other words, given that $Cov\,(\Delta n_t, lp_t)$ has fallen, it must be composed of a drop in $Cov\,(h_t, lp_t)$ and/or a rise in $Cov\,(f_t, lp_t)$. Using different data sources on job flows, I corroborate these conjectures in Figure 2.2.[16]

From the above findings, it is clear that the sudden vanishing procyclicality of productivity around the mid-1980s in the U.S. is not simply an artefact of a particular dataset, or a specific statistical de-trending process, or the choice of the measure of productivity or labour input. Having established the empirical robustness of the so-called *productivity puzzle*, next I investigate its potential cause.

2.3 Drop in Employment Adjustment Cost

Procyclicality of measured productivity in the U.S. after World War II was traditionally explained through *labour hoarding* by firms facing costly hiring and firing of workers. So a natural candidate for explaining the vanishing procyclicality of productivity is a fall in employment adjustment cost. However, whether there has indeed been less factor hoarding after the mid-1980s remains an empirical question. In Section 2.3.1, I study the cyclical properties of factor utilization rate, which is a proxy measure for factor hoarding, and establish that factor hoarding has lost its importance in the post-1980s U.S. In Section 2.3.2, I study the response of the aggregate U.S. economy to technology and demand shocks in a structural VAR set-up. The changes in these responses between the

[14] There is a difference in the levels of the correlations between the two alternative measures of productivity — while TFP has remained procyclical even after the drop, average labour productivity has become countercyclical with hours worked, and acyclical with output. The current thesis is not concerned with these level differences, but the sudden drop around the mid-1980s.

[15] For the complete set of robustness checks for the choice of filters, see Panel A of Table A.1 in the Appendix. Using KLEMS data, provided by Jorgenson, Ho, and Samuels (2012), Panel B of the table shows that the drop in labour productivity correlations is also robust to considering annual data for the aggregate U.S. economy. Data for the non-farm business sector (not shown here) also produce the same correlation pattern. Findings are also robust to using employment as the measure of labour input instead of total hours worked. Murali (2018) shows that the fall in productivity correlations is also robust to using empirical mode decomposition method of calculating the moments at different cyclical frequencies.

[16] It is difficult to obtain data on economy-wide job destruction before the late 1970s. However, economy-wide job vacancy rate can be obtained using the monthly Help Wanted Index (HWI) from the Job Openings and Labor Turnover Survey (JOLTS) starting from 1951. Also, for the manufacturing sector, Davis, Faberman, and Haltiwanger (2006) have collected quarterly data on both job creation and destruction rates starting from 1947.

pre- and post-1984 periods further confirm the hypothesis that firms have resorted to less labour hoarding in recent decades.

2.3.1 Vanishing Procyclicality of Factor Utilization Rate

Commonly used measures of productivity, like labour productivity and TFP, contain an implicit component of factor utilization rate that can itself have cyclical correlations with output and labour input measures like employment and total hours worked. For example, if labour is utilized at a higher rate (by increasing labour effort) during economic booms than during recessions then measured labour productivity will be more procyclical. This can be understood by simply studying a production function with effective labour input, $Y = AE^{\alpha_1}N^{\alpha_2}$, where Y is the value-added, E is the effort or utilization rate of each worker N, and A is the utilization-adjusted productivity component. Average labour productivity is defined as $\frac{Y}{N} = AE^{\alpha_1}N^{\alpha_2-1}$, which is strictly increasing in E and weakly decreasing in N so long as $\alpha_1 > 0$ and $\alpha_2 \leq 1$. In an economic downturn, when firms want to reduce the effective labour input, $E^{\alpha_1}N^{\alpha_2}$, they face the option of either reducing the number of workers N, or decreasing the utilization rate E. When it is costly to adjust employment, firms mostly change effort. As an extreme example, when N is fixed over the business cycle due to costly adjustment, all change in labour productivity is explained by changes in effort. Thus, as firms increase E during booms and decrease it in recessions, labour productivity remains perfectly procyclical since measured productivity is increasing in E. As the cost of adjusting N falls, firms can now rely on reducing N in recessions, thereby boosting labour productivity during economic downturns since productivity is decreasing in N. Thus, lower hiring and firing cost makes measured productivity less procyclical.

This argument of procyclical labour utilization was used to justify the procyclicality of TFP in the post-War U.S. economy. However, it remains to be established whether the drop in cyclical productivity correlations was driven by less factor hoarding or more countercyclical utilization-adjusted productivity. Using hours per worker as a proxy that is proportional to unobserved changes in both labour effort and capital utilization, Basu, Fernald, and Kimball (2001) generated a composite factor utilization rate series and a utilization-adjusted TFP series. Studying the cyclical property of those series in Table 2.1, one can safely conclude that the drop in cyclical correlations of measured productivity is driven by the factor utilization component of TFP, and not the utilization-adjusted 'true' productivity component. As discussed above, factor utilization can become less procyclical if factor adjustment along the extensive margin over the business cycle becomes more pervasive in comparison to changes in unobserved labour effort and work-week of capital.

Notwithstanding the fall in procyclicality of factor utilization rate, utilization-adjusted TFP has historically been and continues to be much less procyclical than factor utilization. Hence, purely in a mechanical variance decomposition sense, if the relative contribution of factor utilization rate falls in the total variability of aggregate TFP, measured productivity will become more countercyclical. Table 2.2 shows that the share of the total variation of TFP that is explained by the more procyclical component of factor utilization rate has diminished sharply in the post-1984 period. Such a shift towards greater relative importance of the extensive margin of factor adjustment can emanate from a drop in the cost of hiring and firing of factors of production, particularly labour.[17]

Falling employment adjustment cost should imply a rise in the volatility of employment relative to those of output and factor utilization. Panels (a) and (b) of Figure 2.3 show the dramatic rise in the volatility of hours and employment relative to that of output exactly at the time of the sudden drop in the productivity correlations. Finally, Panel (c) of Figure 2.3 shows how the relative importance of employment (the extensive margin of labour adjustment) vis-à-vis the intensive margin of factor utilization has progressively increased from around the same time. This rise in the relative volatilities of measured labour inputs happened immediately after the onset of the so-called Great Moderation when the absolute volatilities of output and labour input fell precipitously in the late 1970s. As is evident from Figure 2.4, even though the volatilities of output, hours and employment follow a similar time trend, the magnitude of reduction in volatility is larger for output than for the labour inputs. This leads to the eventual increase in the volatility of labour input relative to that of output.

To summarize, the vanishing procyclicality and reduced volatility of factor utilization rate over the business cycle, induced by a drop in employment adjustment cost, can not only explain the fall in measured productivity correlations but also the rise in relative volatility of employment.

2.3.2 Changes in Response to Technology and Demand Shocks

Structural changes in the labour market that make hiring and firing of workers easier for firms should have implications for how the economy responds to different types of shocks. In this section, I focus on such changes in the response of the aggregate U.S. economy to technology and demand shocks between the pre and post-1984 periods. I perform a time-varying structural vector autoregression (SVAR) with labour productivity growth and per capita hours, as in Galí and Gambetti (2009).[18] There are two main advantages of this specification: first, it allows one to control for low-frequency movements in per capita hours without having to extract the cyclical component of hours through any form of ad hoc time-series filtering, and second, it allows one to know the complete dynamics of the impulse responses over the years so that it can be pinpointed as to exactly when the responses began to change.[19] Since the current thesis focusses on documenting the changes in the impulse responses during the mid-1980s, this method of time-varying SVAR is the most suitable for the purpose. For a detailed discussion on the choice of SVAR specification and identification of the shocks, refer to Appendix A.3.

[17]See Appendix Table A.4 for the robustness of these results using an alternative measure of capacity utilization rate published by the Federal Reserve Board.

[18]Chang and Hong (2006) criticize the use of labour productivity as a measure of productivity. They argue that using labour productivity instead of TFP mislabels changes in input mix (i.e., permanent changes in the capital-labour ratio) as technology shocks. Hence, as a robustness check, I perform the same SVAR replacing labour productivity with TFP.

[19]In Figures A.4 and A.6, I show the dynamics of the impulse responses by year. This confirms the choice of 1984 as the approximate year of the structural change, although no rigorous structural change test was performed.

Response to technology shock. Panels (a), (b) and (c) of Figure 2.6 respectively show the impulse responses of per capita hours, per capita output and labour productivity to a positive technology shock separately for the pre-1983 (solid blue lines) and post-1984 (dashed red lines) periods. Of these, the only statistically significant difference between the two sub-periods is the change in the impulse response of hours, as shown in Appendix Figure A.5. While I find that per capita hours respond negatively on impact to a positive technology shock throughout the post-War era, the negative response is much less intense and barely different from zero in the post-1984 period.[20] A common explanation provided for the diminished response of hours to technology shocks is that the monetary policy conducted by the Federal Reserve became more accommodative of technology shocks to the economy. But what is most relevant in the context of the productivity puzzle is that the muted negative response of hours to a positive technology shock increases the productivity correlation with labour input. This acts as a counterforce to the vanishing procyclicality of productivity.

Response to demand shock. In response to a positive demand shock, hours increased by roughly the same amount in the pre- and post-1984 periods, while the positive response of output on impact was drastically muted after the mid-1980s. Since average labour productivity is nothing but output per hour worked, the reduced response of output and a near-identical response of hours implies a muted response of labour productivity to a demand shock.[21] A qualitatively similar result is obtained when estimating the SVAR with TFP growth instead of labour productivity growth. While the impulse response of TFP remains positive albeit the reduction in magnitude (see Panel (d) of Appendix Figure A.7), the on-impact response of labour productivity turns negative after 1984 (see Panel (f) of Figure 2.6). The muted response of productivity and an unchanged response of hours to a demand shock imply that conditional on a demand shock, the correlation of productivity with hours must drop in the post-1984 era. This is shown in Figure 2.7.

The reduction in productivity correlation conditional on a demand shock proves that it is not the case of changing composition of shocks to the U.S. economy that have induced the sudden fall in unconditional productivity correlation, rather there must have been deeper structural changes in the economy that caused firms to change output by a smaller magnitude when hit with the same demand shock. An example of such a structural change is the decline in the labour adjustment cost. Given a positive demand shock, when employment adjustment is less costly, a firm does not increase the intensive margin of effort as much, which causes output and productivity to not rise as much for a given increase in employment and hours worked. This decreases the correlation of productivity with measured labour input.

[20] Contrary to my findings, Galí and Gambetti (2009) did not find a starkly muted response of hours conditional on a technology shock in the post-1984 period. This difference emanates from extending the post-1984 period with more recent years of data — while they used data till 2005, my dataset extends till the fourth quarter of 2017. When TFP is used as the measure of productivity instead of labour productivity, a similar difference between the two sub-periods in the initial response of hours emerges (see Panel (a) of Appendix Figures A.7 and A.8).

[21] For statistical significance of the differences in the impulse responses between the two sub-periods, refer to Panels (d), (e) and (f) of Appendix Figure A.5.

In summary, while the change in response to demand shocks predicts decreasing procyclicality of productivity, the change in response to technology shocks predicts just the opposite. How all these opposing forces combine to generate the change in the unconditional correlation of productivity with output and hours will be studied in Chapter 3. Nevertheless, the crucial finding is that the unconditional correlation of productivity with output and hours fell not because of the rising importance of technology shocks vis-à-vis demand shocks (as claimed by Barnichon (2010)) but because of a drop in the correlations *conditional* on demand shock, thereby pointing towards structural changes in the economy that made factor hoarding less relevant.

2.4 What Caused Employment Adjustment Cost to Drop

The reduced importance of factor utilization rate in measured productivity and the change in productivity correlation conditional on demand shocks establish that higher dependence on hiring and firing of workers instead of the intensive margin of effort adjustment has caused the procyclicality of productivity to fall so drastically.[22] However, what observable structural change in the labour market can bring about such a sudden drop in employment adjustment cost remains an open question so far.

One such possible cause of increasing employment turnover is the rise in online job-search platforms, which reduces the hiring cost by making it much easier to match workers and jobs. Moreover, the improved efficiency of online matching between specific worker and job types could also mean that firms need to terminate fewer workers who do not fit well with the job, thereby reducing the firing cost for firms. However, this is unlikely to have triggered the switch in the productivity correlations in the mid-1980s because internet recruitment service providers did not begin their journey until the mid-1990s.

The increased use of temporary workers is another likely reason for the reduction in employment adjustment cost. Jalón, Sosvilla-Rivero, and Herce (2017) argue that the countercyclicality of labour productivity in Spain was driven by the 1984 legislative reform that increased the importance of temporary workers in the Spanish economy. Daruich, Addario, and Saggio (2017) also study the implications of a similar 2001-reform of lifting constraints on the employment of temporary contract workers in Italy. For the U.S. it is difficult to ascertain the role of temporary workers in the increased flexibility of labour markets due to lack of suitable data that dates back long enough, e.g., employment data for the temporary help services industry from the Current Employment Statistics (CES) database of BLS dates back only till 1990. Although Carey and Hazelbaker (1986) show that employment growth in the temporary help industry increased sharply immediately after the 1982 recession, which lines up well with the timing of the switch in labour productivity correlations, Schreft and Singh (2003) show that temporary and part-time hiring and overtime — collectively

[22] There could have been a similar drop in the cost of adjusting capital stock along the extensive margin along with or instead of a drop in the cost of hiring and firing workers. However, the proxy used here to measure the factor utilization rate is hours per worker, as advocated by Fernald (2014). Since it is empirically impossible to make a distinction between labour utilization (effort) and capital utilization rates, and because hours per worker is arguably a more direct proxy for labour effort, I will stick to only employment adjustment cost in this thesis.

known as 'just-in-time hiring' — has gained in importance only since the 1991 recession in the U.S. However, for the U.S., I study the time series of the share of part-time workers (see Figure 2.5) and do not find any noticeable upsurge, if not an actual plateauing, in the share of part-time workers around the mid-1980s.

Galí and van Rens (2017) claim that the main driver of falling labour market frictions in the U.S. labour market was the drop in job separation rate. They argue that because of a substantial drop in the gross job destruction rate, firms need to hire much less new workers to maintain the level of employment. This reduced hiring activity implies lower cost of employment adjustment in equilibrium, thereby leading to more countercyclical productivity. While this channel of reduction in employment adjustment cost is certainly feasible for the U.S., a cursory glance at the international evidence from twelve other OECD countries, presented in Table 2.3, essentially refutes the claim that change in job separation rate is a significant determinant of changes in productivity correlations.[23]

2.4.1 Decline of Unions

An empirically identifiable labour market change that occurred in the U.S. almost around the same time as the change in the productivity correlations is the decrease in size and influence of labour unions. In Figure 2.8 we see that union membership among working individuals (both in terms of rates and absolute numbers) was rising in the U.S. until the early 1970s, after which it remained flat for a decade (with falling rates for the private industries and increasing rates for the public sector), and started falling sharply since the early 1980s.[24] To emphasize how dramatic the de-unionization event around the 1980s was, one can compare the growth rates in union density for 30 years before and after 1980. In the three decades preceding 1980, unionization rate remained almost constant, while between 1980 and 2010 it fell by roughly 50% in aggregate and by 67% in the private sector. One concern about de-unionization being the main driving force behind falling procyclicality of productivity is that union rates were already quite low in the U.S. even before 1980, roughly 20% of the workforce and so falling union rate should not matter much. Taschereau-Dumouchel (2017) argues that it is not so much the fraction of union-jobs but the presence of the political threat of unionization that matters for labour market outcomes.

Farber and Western (2002) argue that this stark reversal of unionization trend in the U.S. was precipitated by a fall in the annual number of union elections — a key channel of recruiting new union members. The number of elections held fell by almost 50%, from about 8000 in 1980 to about

[23]Of the countries considered, only Ireland experienced a notable decrease in the job separation rate along with decreasing cyclical correlation of labour productivity. Nevertheless, Ireland also experienced a 21% drop in union density, and hence the exact source of its vanishing procyclicality of productivity cannot be determined easily.

[24]Consistent data on union density is available separately for the private sector only from 1973 onwards (see Hirsch and Macpherson (2003)). Although unionization rate started falling from the early 1970s in the private sector, the de-unionization process accelerated from 1980: the average annual rate of decline in private-sector union density was 2.4% between 1974 and 1979 compared to 6.6% between 1980 and 1985. Troy and Sheflin (1985) present data on private-sector union density in the U.S. between 1929 and 1972, and they also find an average annual de-unionization rate of only 1.1% between 1950 and 1972. Therefore, it can be concluded that the decline of unions even in the private sector had a sharp acceleration starting from the early 1980s.

4400 in 1985. The unfavourable political climate was precipitated by President Reagan's strong stand against the air-traffic controllers' strike of 1981,[25] and the much-publicized appointment of the Reagan Labor Board in 1983. A change in the political climate regarding labour unions can also mean that changes in union density might be an underestimate of the change in the real bargaining power of unions. While it is difficult to directly measure the power of unions, one good proxy is to look at the number of work stoppages, which are usually organized by unions. From Figure 2.10, one can see that large-scale work stoppages dropped by almost 90% of its pre-1980 level quite suddenly within a couple of years. Thus, although the decline in union membership from the early 1980s was a somewhat gradual process which might seem inconsistent as an explanation for the strikingly rapid decline in the productivity correlations, union power seems to have declined more promptly. Moreover, in Figure 2.9, I show that the timing and speed of de-unionization match well with that of the productivity puzzle.

2.4.2 De-unionization: International Evidence

The era of deregulation that began in the United States from the early 1980s had its parallel in other parts of the world.[26] It is interesting to note as anecdotal evidence that the United Kingdom, which underwent a similar de-regulation episode under Margaret Thatcher, experienced both de-unionization and drop in procyclicality of productivity. On the other hand, countries like Canada, for which this drop in unionization is conspicuously absent (see Riddell (1993)), did not undergo a drop in cyclical correlations of productivity. In Table 2.3, I show that in most of the developed world, de-unionization is strongly predictive of the loss in productivity procyclicality and a rise in the volatility of employment relative to that of output.[27] This evidence is consistent in spirit with Gnocchi and Pappa (2009), who find that union coverage is one labour market rigidity that most significantly affects business cycle statistics in a sample of 20 OECD countries.

The evidence that de-unionization did not happen in all industrial countries highlights another important aspect of the phenomenon. It is natural to link de-unionization with other labour market changes happening at the same time as its potential cause. For example, Acemoglu, Aghion, and Violante (2001) and Dinlersoz and Greenwood (2016) argue that skill-biased technological change can explain de-unionization in the U.S., and Açikgöz and Kaymak (2014) show that roughly 40% of the drop in unionization rates in the U.S. can be explained by the rise in the skill premium in wages. Further, Foll and Hartmann (2019) argues that routine task-biased technical change is not only the driving force behind job market polarization but also de-unionization. However, since skill-biased or routine-biased technological change happened in most developed economies, the sudden

[25]On August 5, 1981, Reagan fired more than 11,000 striking air traffic controllers who had ignored his order to return to work. This sweeping mass firing of federal employees sent a strong message to American business leaders that they can hire and fire their workers much more easily.

[26]De-regulation in the labour market can not only mean the decline in union power but also relaxation of employment protection legislation (EPL). However, as shown in Table 2.3, I do not find any consistent pattern across OECD countries that less stringent EPL translated into lower procyclicality of productivity.

[27]Figure 2.18 shows the regression between changes in union density and the changes in labour productivity correlation in a cross-section of 13 developed economies.

trend reversal in union density in only a handful of these countries like the U.S. is likely to be mostly driven by political factors. Moreover, insofar as one believes that skill-biased technological change from the 1980s was driven by IT capital use (due to high capital-skill complementarity in the production process, e.g., as highlighted in Krusell, Ohanian, Rios-Rull, and Violante (2000)), one should find a significant correlation across industries between the rising share of IT capital and falling productivity correlations. This is however not the case, as pointed out by Wang (2014). Therefore, while it could be the case that relatively slow-moving technological changes impacting the labour market had some role to play in the de-unionization process, the episode of rapid fall in union power from the early 1980s is most likely to have been precipitated by political factors that are exogenous to labour market conditions. It is in this sense of exogeneity that the impact of de-unionization on the falling procyclicality of productivity (as shown in the next section using cross-sectional variation from U.S. states and industries) can be thought as a *causal* channel.

2.4.3 De-unionization: Evidence from U.S. States and Industries

Having established that in the aggregate U.S. economy, de-unionization since the early 1980s is consistent, in terms of timing, with falling procyclicality of labour productivity, and rising volatility of employment relative to that of output and factor utilization, I now use sectoral variation across U.S. states and industries to see if a larger magnitude of de-unionization is indeed correlated with a greater reduction in labour productivity correlation. In particular, I run the following cross-sectional regression:

$$\Delta Corr\left(lp_i, h_i\right) = \alpha + \beta \Delta \ln \left(Union\ Density\right)_i + \varepsilon_i, \tag{2.1}$$

where lp_i and h_i are the cyclical components of labour productivity and hours in industry or state i, and the time-difference Δ denotes the difference between the pre- and post-1984 average values.[28]

Figures 2.11 and 2.15 show a significant positive relationship between the degree of de-unionization and the magnitude of the drop in productivity correlations across 17 U.S. industries and 51 U.S. states respectively. To avoid the result being driven by small industries or states, I weight the observations with the average employment level in each industry or state.

One concern while using cross-sectional variation in de-unionization in the above regression is that union densities might not have fallen substantially within individual industries or states. In other words, the fall in aggregate unionization rate might have been driven by employment shifts towards less unionized sectors and regions, rather than de-unionization within them. This would be problematic for identifying the slope coefficient in regression (2.1), because of lower cross-sectional variation in changes in union densities. However, a simple within-between decomposition[29] of aggregate de-unionization shows that nearly 88% of the fall in aggregate union density happened

[28] I do not have data with simultaneous variation across states and industries, and so a single regression exploiting that *industry* $\times$ *state* variation could not be performed.

[29] Total change in union density, Δu = Within-i change, $\sum_{i=1}^{17} \bar{e}_i \Delta u_i$ + Between-i change, $\sum_{i=1}^{17} \bar{u}_i \Delta e_i$, where $\bar{e}_i$ is the average employment share and $\bar{u}_i$ is the average union density in industry or state i.

within the 17 industries considered here between 1983 and 1991. Similarly, 91% of the total fall in the unionization rates in the U.S. between the pre- and post-1984 periods took place within the states, and not through employment shifts towards less unionized states.

For the state-level regression, there is an additional concern that in recent years many U.S. states have adopted *right-to-work* legislation promoting their "pro-business" outlook, thereby rendering the labour unions a lot less powerful in those states. In that case, a decline in union density in these right-to-work states should barely matter for explaining the drop in productivity correlations. In Appendix Figure 2.16, I show this is indeed the case, with only the so-called non-right-to-work states driving the positive relationship between de-unionization and drop in productivity correlation. This finding of right-to-work law interacting with union power to determine productivity through changes in management practices resonates well with U.S. plant-level findings by Bloom, Brynjolfsson, Foster, Jarmin, Patnaik, Saporta-Eksten, and Van Reenen (2019).

One alternative identification strategy to the one considered above, is to perform a difference-in-difference estimation à la Card (1992). In that strategy, one assumes that the intensity of the de-unionization event is larger in industries with a higher initial proportion of unionized workers. Thus, instead of regressing the change in the productivity correlation on the change in the union density, one simply regresses it on the pre-1984 level of union density:

$$\Delta Corr\,(lp_i, h_i) = \alpha + \beta \ln\left(Union\ Density\right)_i^{pre-1984} + \varepsilon_i \tag{2.2}$$

This method of identification also corroborates my finding that union density had a significant role to play in the vanishing procyclicality of labour productivity and the rising volatility of employment changes relative to that of output.[30]

The above cross-sectional evidence supports my claim that de-unionization is positively correlated with vanishing procyclicality of labour productivity, but it is still left to show that industries with a greater fall in productivity correlation experienced a larger increase in the volatility of employment (that is, the extensive margin of labour adjustment) relative to that of output and utilization (that is, the intensive margin of factor adjustment). Since utilization data is not available at the industry level, I use hours per worker as a proxy for the measure of the intensive margin of labour adjustment. From Figure 2.17 one can find a statistically significant negative relationship between the change in labour productivity correlation on one hand and the relative volatility of employment on the other, across U.S. industries.[31]

The negative correlation patterns in Panels (a) and (b) of Figure 2.17 are similar, but there is a subtle difference between the two scatter plots. A lot of the industries experienced a rise in the volatility of employment relative to that of output, while very few experienced a similar rise in the volatility of employment relative to that of hours per worker. This finding that the extensive margin of employment adjustment became less volatile relative to the intensive margin of changing hours

[30] For details on the difference-in-difference strategy and the results based on that, refer to Appendix A.2.

[31] As more direct evidence of the union channel, Figure 2.12 shows the negative correlation between larger declines in the union density and the rising relative volatility of employment across U.S. industries.

per worker in almost all industries is an apparent aberration from what one would expect under falling employment adjustment cost.[32] This is particularly puzzling in light of the evidence in Panel (c) of Figure 2.3 that employment became increasingly more volatile relative to factor utilization since the 1980s. To understand why the dynamics of hours per worker and factor utilization might vary, it is instructive to study the industry-level differences in the elasticity of output to hours per worker. Basu, Fernald, and Kimball (2001) find that the responsiveness of output to hours per worker is vastly different across industry-groups, e.g., the non-durables manufacturing sector is roughly 60% more responsive than the durables manufacturing industries, and more than thrice as responsive as the service sector. The declining share of manufacturing in the U.S. can, therefore, explain the larger decline in the volatility of factor utilization at the aggregate level than within individual industries.

2.5 Conclusion

Why did productivity suddenly start becoming less procyclical during the mid-1980s? This is the research question that this chapter tries to answer. A decomposition of measured productivity into factor utilization and a utilization-adjusted productivity component reveals that the fall in productivity correlations is driven by a lower dependence on factor hoarding. Changes in responses of the U.S. economy to demand shocks also point towards some structural change that made adjusting factors along the extensive margin less costly, thereby leading to firms relying less on factor hoarding. I then identify a rapid decline in union power since the early-1980s as the key structural change in the labour market that made hiring and firing of workers easier relative to effort adjustment through labour hoarding. Using cross-sectional evidence from U.S. industries and states, and also international evidence from OECD countries, it is argued that more intense de-unionization is associated with a deeper fall in the cyclical productivity correlations. Moreover, lower dependence on labour hoarding in the face of higher labour market flexibility is shown to imply rising relative volatility of employment, and this rise is also correlated with the fall in union density and productivity correlations across U.S. industries. Understanding the role of labour market institutions like unions in influencing business cycle properties of aggregate macroeconomic variables is the broad contribution of this chapter.

[32] The post-1984 change in the volatility of employment relative to that of hours per worker is not very robust to the choice of different datasets and time-series filters (see Table A.3 in Appendix). Regardless of the filter used, the volatility of employment relative to hours per worker did not show any stark upward trend after the mid-1980s.

2.6 Tables

Table 2.1: Reduction in Procyclicality of Factor Utilization Rate

Variable & Filter Choice	**With Output** Pre 1983	Post 1984	Change	**With Hours** Pre 1983	Post 1984	Change
Panel A: Quarterly Growth Rate						
TFP	0.87	0.70	- 0.17	0.35	0.10	- 0.25
Factor Utilization Rate	0.73	0.49	- 0.24	0.67	0.52	- 0.15
Utilization-Adjusted TFP	0.10	0.25	+0.15	-0.40	-0.32	+0.08
Panel B: Annual Growth Rate						
TFP	0.88	0.69	- 0.19	0.49	0.29	- 0.20
Factor Utilization Rate	0.87	0.62	- 0.25	0.75	0.61	- 0.15
Utilization-Adjusted TFP	-0.10	0.04	+0.14	-0.45	-0.34	+0.11

Note: Data on quarterly and annual growth rates of TFP, factor utilization rate, utilization-adjusted TFP, output and hours worked for the U.S. business sector are sourced from Fernald (2014). Since Fernald (2014) only provides the growth rates of the three variables, robustness to other de-trending methods cannot be established.

Table 2.2: Reduction in Variance of Factor Utilization Rate

Variable & Filter Choice	**Variances** 1948-1983	1984-2017
Panel A: Quarterly Growth Rate		
Total Factor Productivity (TFP)	17.55 (100%)	5.89 (100%)
Factor Utilization Rate	11.67 (66.5%)	1.64 (27.8%)
Utilization-Adjusted TFP	5.88 (33.5%)	4.25 (72.2%)
Panel B: Annual Growth Rate		
Total Factor Productivity (TFP)	4.56 (100%)	1.50 (100%)
Factor Utilization Rate	3.79 (83.1%)	0.79 (52.7%)
Utilization-Adjusted TFP	0.77 (16.9%)	0.71 (47.3%)

Note: Data on quarterly and annual growth rates of TFP, factor utilization rate and utilization-adjusted TFP are sourced from Fernald (2014). Since Fernald (2014) only provides the growth rates of the three variables, robustness to other de-trending methods cannot be established. Percentages in parentheses refer to the share of total variance of TFP that is explained by each component. While calculating the variance of the components, the covariance term was equally split, e.g., $[Var\,(\text{Factor Utilization Rate}) + Cov\,(\text{Factor Utilization Rate, Utilization-Adjusted TFP})]$ is the variance of factor utilization rate, and the variance of utilization-adjusted TFP is given by $Var\,(\text{Utilization-Adjusted TFP})$ $+Cov\,(\text{Factor Utilization Rate, Utilization-Adjusted TFP})$.

Table 2.3: Labour Market Statistics from OECD Countries

Country	ΔCorrelation of Productivity With Output	ΔCorrelation of Productivity With Hours	$\Delta \frac{\text{S.D.(Employment)}}{\text{S.D.(Output)}}$	Labour Market Structure ΔUnion	Labour Market Structure ΔSeparation Rate	Labour Market Structure ΔEPRC
France	-0.13	0.17	26%	-54%	0%	1%
U.S.A.	**-0.54**	**-0.62**	**32%**	**-49%**	**-24%**	**0%**
Australia	-0.44	-0.48	73%	-37%	4%	21%
Austria	-0.21	-0.16	-16%	-32%	No data	-11%
U.K.	-0.39	-0.46	41%	-28%	11%	16%
Spain	-1.37	-0.74	317%	-24%	-1%	-34%
Germany	-0.04	-0.52	-10%	-24%	41%	8%
Ireland	-0.44	-0.21	44%	-21%	-44%	-2%
Italy	-0.09	-0.16	71%	-4%	11%	0%
Norway	-0.35	-0.12	47%	-3%	47%	0%
Canada	0.01	0.09	-22%	2%	9%	0%
Sweden	0.01	-0.03	59%	10%	84%	-7%
Finland	-0.25	0.21	-9%	36%	No data	-22%

Note: Countries are arranged in ascending order of union density changes. All changes are between the post and pre-1984 periods. Productivity is defined as real GDP per hour worked. De-trending of variables has been done using the HP-filter. Quarterly data on output and hours between 1960 and 2010 for all countries (except Spain) are taken from *OECD Economic Outlook Database*, collected by Ohanian and Raffo (2012). Annual data for Spain between 1950 and 2017 is sourced from the *Conference Board Total Economy Database*. Union density data are sourced from *OECD Annual Trade Union Density Dataset*. Since internationally comparable data on job flows are not available before 1980s, changes in job separation rate are calculated as the difference between the average rate between 2002 through 2007, and that between 1985 through 1990, as reported in Elsby, Hobijn, and Sahin (2015). Employment protection is the EPRC index from the OECD from 1985 to 2013. The index is very persistent over time, so changing the end year of the sample would make very little difference.

2.7 Figures

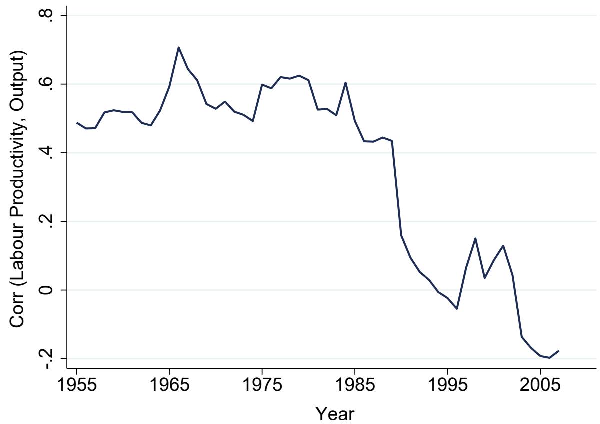

(a) Corr (Labour Productivity, Output)

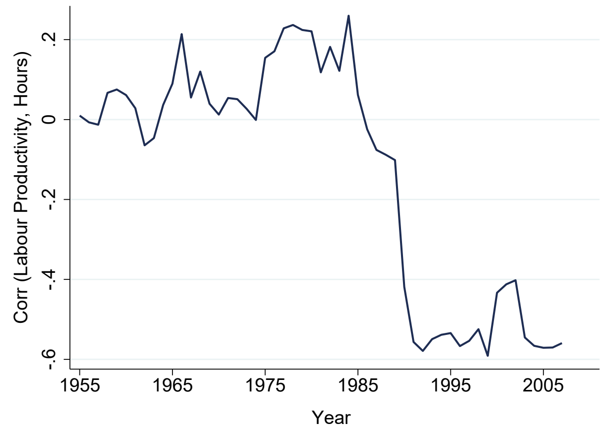

(b) Corr (Labour Productivity, Hours)

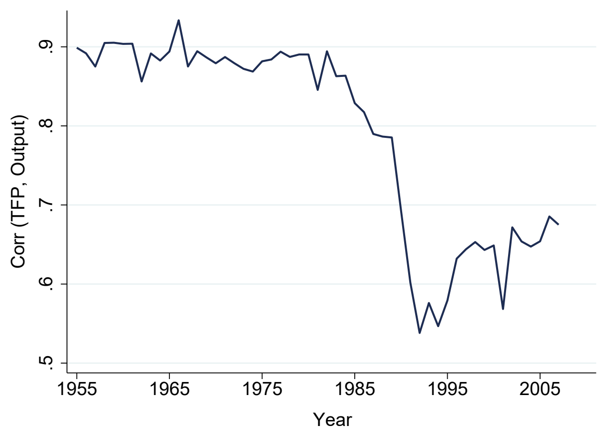

(c) Corr (TFP, Output)

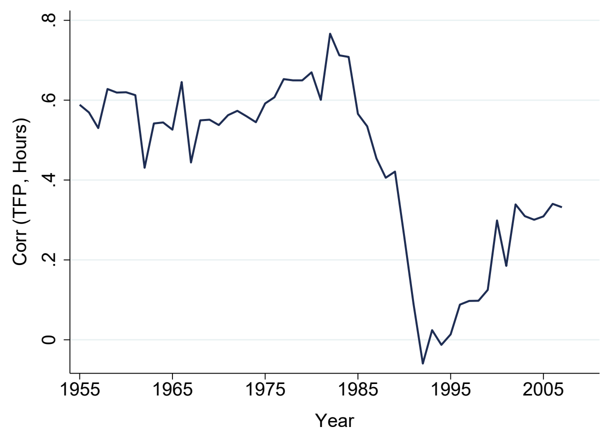

(d) Corr (TFP, Hours)

Figure 2.1: Vanishing Procyclicality of Productivity in the United States

Note: Output, hours and average labour productivity (output per hour worked) data for Panels (a) and (b) are sourced from the *Labor Productivity and Costs* quarterly dataset published by the *Bureau of Labor Statistics* for the U.S. business sector. Relevant data for Panels (c) and (d) are sourced from Fernald (2014), as modified by Ramey (2016). The measure of TFP is not adjusted for factor utilization. The Baxter and King (1999) bandpass filter between 6 and 32 quarters is used to filter all the variables. A centred rolling window of 15 years is used to calculate the correlations. Findings are robust to alternative choice of filters and window-sizes.

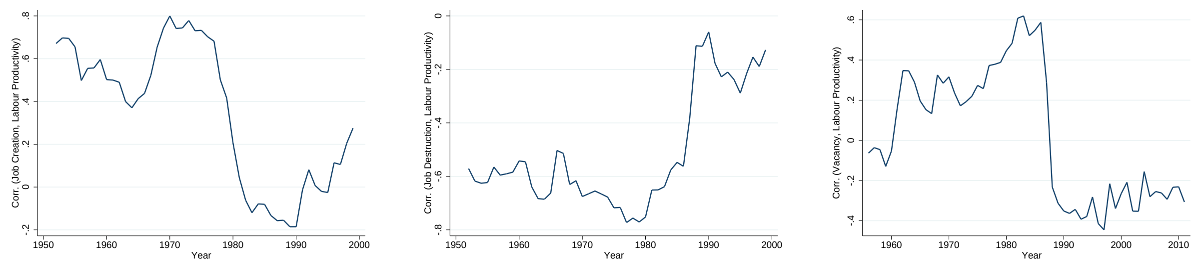

(a) Job Creation (b) Job Destruction (c) Help Wanted Index

Figure 2.2: Cyclical Correlation of Labour Productivity with Job Flows

Note: Panels (a) and (b) correspond to the U.S. manufacturing sector (data from Davis, Faberman, and Haltiwanger (2006)), while Panel (c) is for the entire U.S. economy (data from the *Job Openings and Labor Turnover Survey*). The Baxter and King (1999) bandpass filter between 6 and 32 quarters is used to filter all the variables. A centred rolling window of 10 years is used to calculate the correlations. Findings are robust to alternative choice of filters and window-sizes.

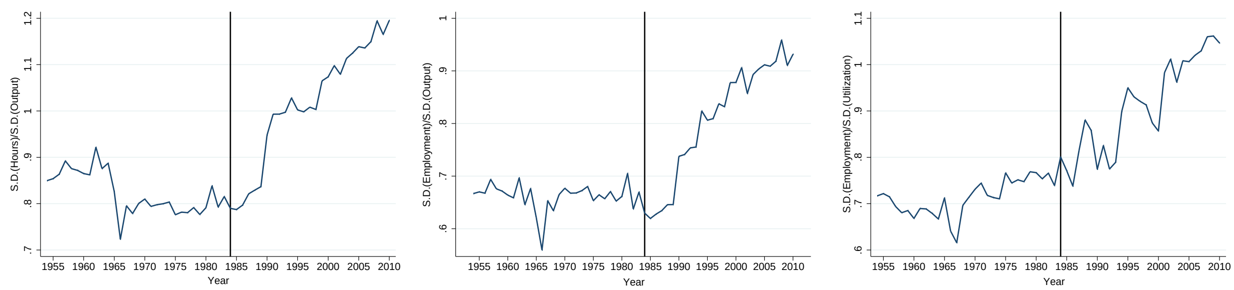

(a) $\frac{\text{s.d.(Hours)}}{\text{s.d.(Output)}}$ (b) $\frac{\text{s.d.(Employment)}}{\text{s.d.(Output)}}$ (c) $\frac{\text{s.d.(Employment)}}{\text{s.d.(Utilization)}}$

Figure 2.3: Relative Volatility of Hours & Employment over the Business Cycle

Note: Data for hours, employment and output is sourced from the BLS-LPC quarterly dataset for the U.S. business sector. Factor utilization data (in quarterly growth rates) is taken from Fernald (2014). The Christiano and Fitzgerald (2003) bandpass filter between 6 and 32 quarters have been used to extract the cyclical component of the variables in Panels (a) and (b), while the annualized quarterly growth rate has been used in Panel (c). A centred rolling window of 15 years is used to calculate the second moments. Findings are robust to alternative choice of filters and window-sizes.

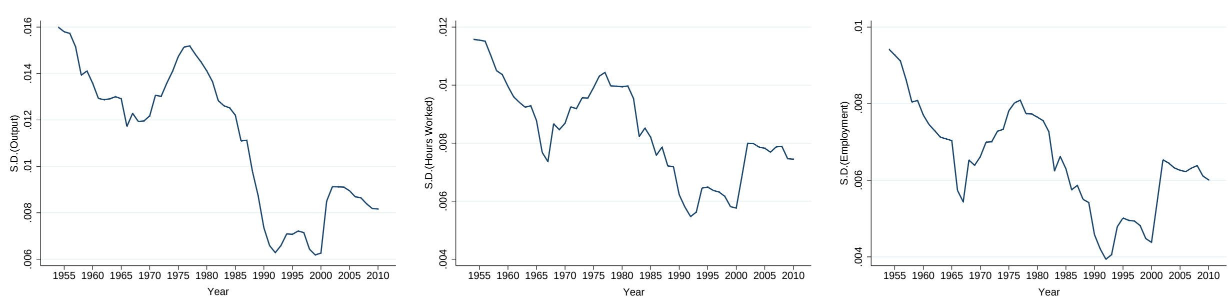

(a) Output (b) Hours (c) Employment

Figure 2.4: Cyclical Volatility of Quarterly Growth Rates of Output, Hours & Employment

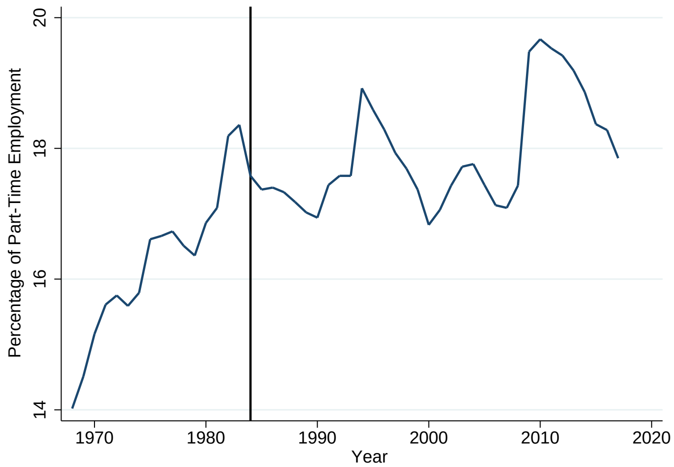

Figure 2.5: Share of Part-time Employment in the U.S. (1968-2017)

Note: Data is sourced from Labor Force Statistics (LFS) of the Current Population Survey (CPS). Part-time employment is defined as less than 35 hours of work per week.

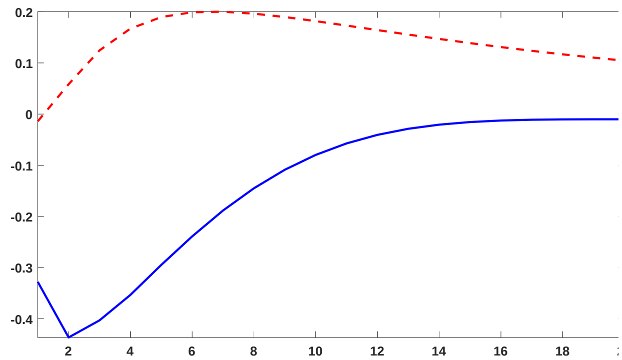

(a) Technology Shock: Hours

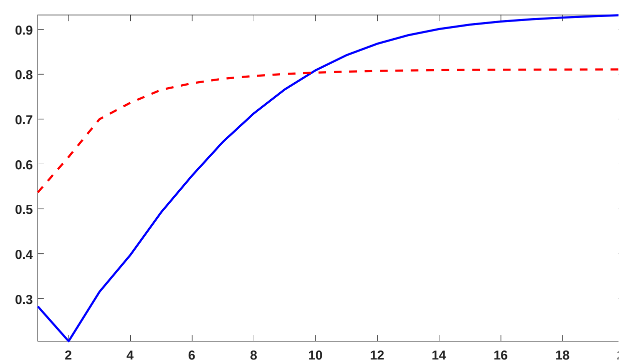

(b) Technology Shock: Output

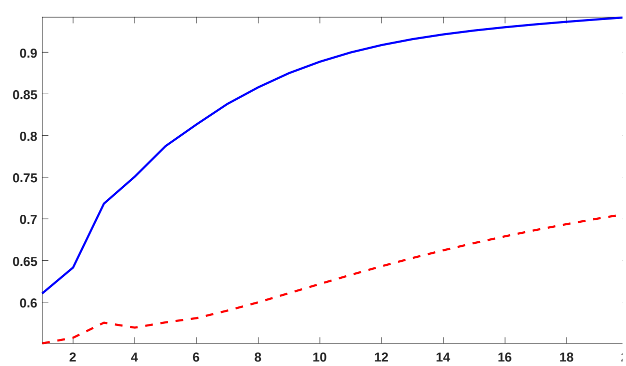

(c) Technology Shock: Labour Productivity

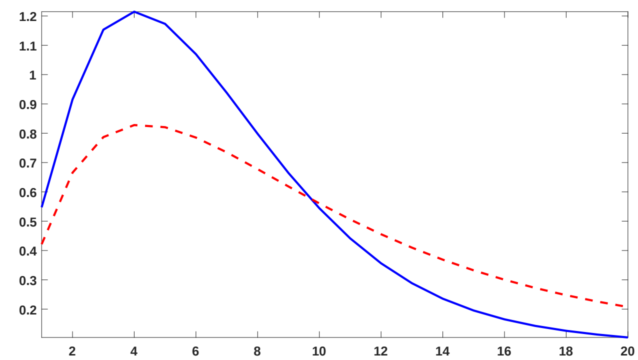

(d) Demand Shock: Hours

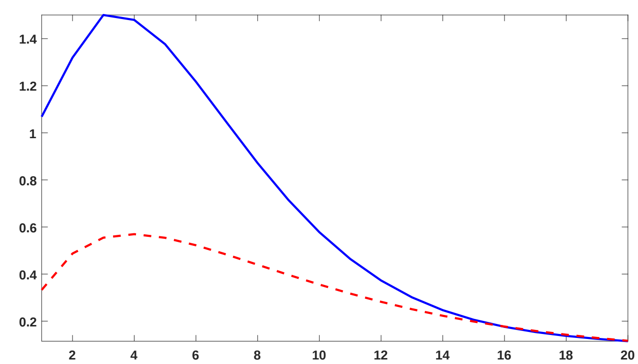

(e) Demand Shock: Output

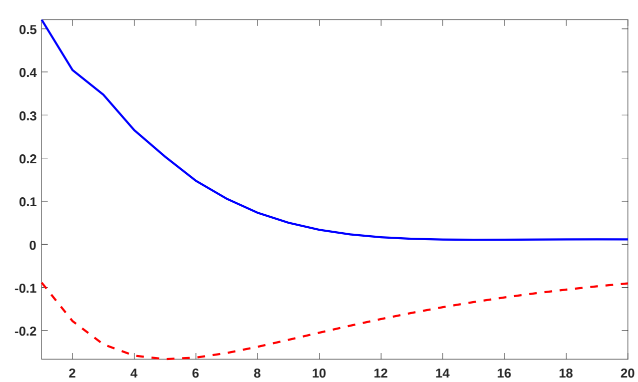

(f) Demand Shock: Labour Productivity

Figure 2.6: Impulse Responses to Technology & Demand Shocks (LP, Hours & Output)

Note: Impulse Response Functions (IRF's) of per-capita hours, per-capita output and average labour productivity from a 2-variable (viz., labour productivity growth and per-capita hours) time-varying long-run SVAR. IRF's for the pre-1984 period (1956-1983) are in blue, and the post-1984 (1984-2017) IRF's are in red dashed lines. Data is sourced from the *Labor Productivity and Costs* (LPC) quarterly dataset for the U.S. business sector, published by the *Bureau of Labor Statistics* (BLS).

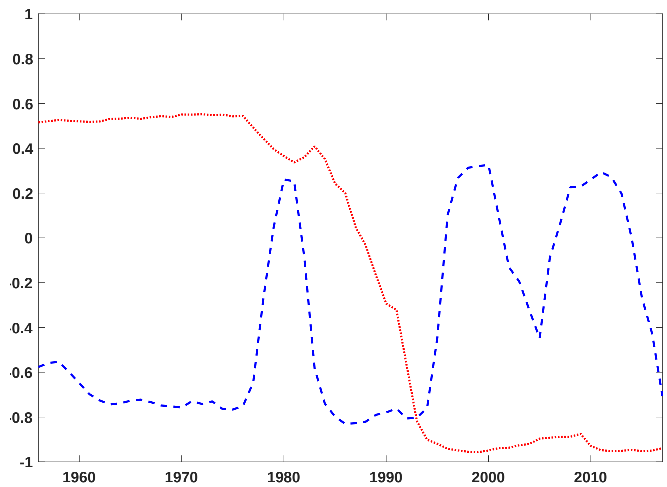

Figure 2.7: Conditional Correlations of Labour Productivity with Hours

Note: Time-varying correlations of per capita hours with labour productivity, conditional on technology shock (blue dashed line) and demand shock (red dotted line).

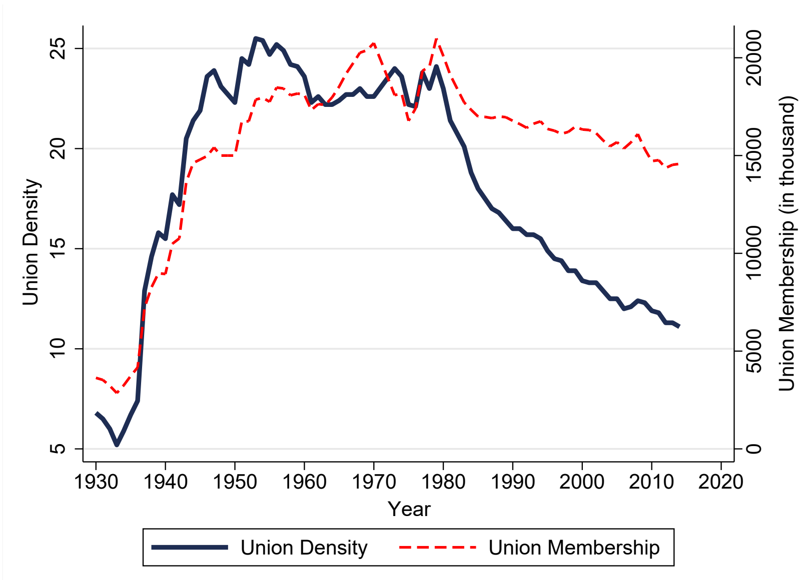

Figure 2.8: Size & Density of Labour Union Membership in the U.S. (1930-2014)

Note: Figures represent the number and percentage of non-agricultural wage and salary employees who are union members. Data before 1977 is sourced from Historical Tables published by the *Bureau of Labor Statistics*. Data between 1977 and 1981 comes from May earnings files, and from 1983 onwards it comes from the Outgoing Rotation Group (ORG) earnings files of the *Current Population Survey* (CPS), collected by Hirsch and Macpherson (2003). Union coverage rates are slightly different from union membership rates but follow a similar time-trend.

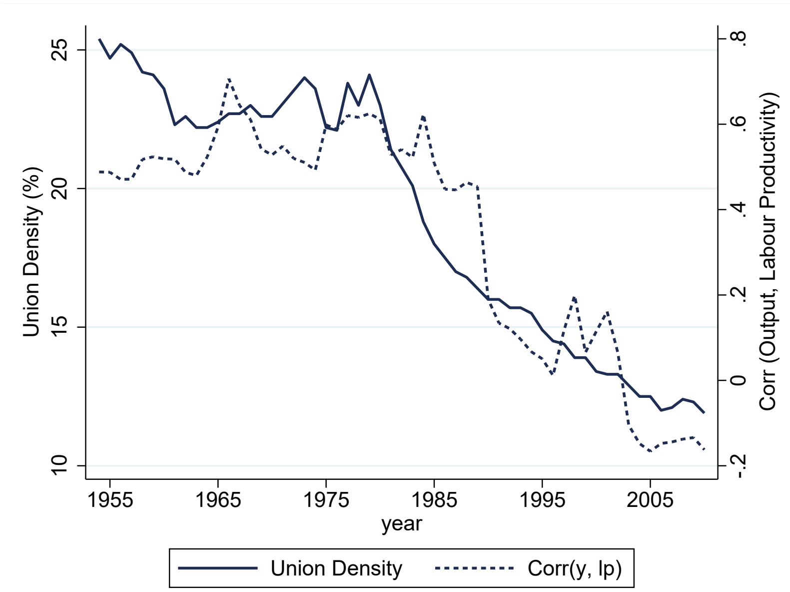

Figure 2.9: De-unionization and Vanishing Procyclicality of Productivity

Note: Union density, sourced from Hirsch and Macpherson (2003) and Historical Tables from the *Bureau of Labor Statistics* (BLS), refers to the percentage of non-farm wage and salary employees who are union members. $Corr(y, lp)$ refers to a 15-year rolling window correlation between the Baxter and King (1999) filtered (between 6 and 32 quarters) business sector output and average labour productivity, sourced from the *Labor Productivity and Costs* quarterly dataset of the BLS.

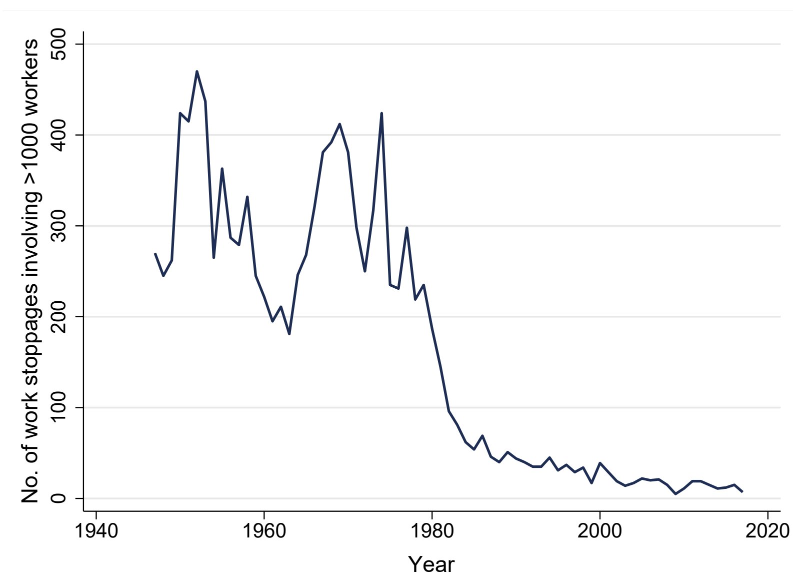

Figure 2.10: Number of Work Stoppages involving 1000 or more workers in the U.S. (1947-2017)

Note: Data is sourced from the Economic News Release of the Bureau of Labor Statistics (BLS).

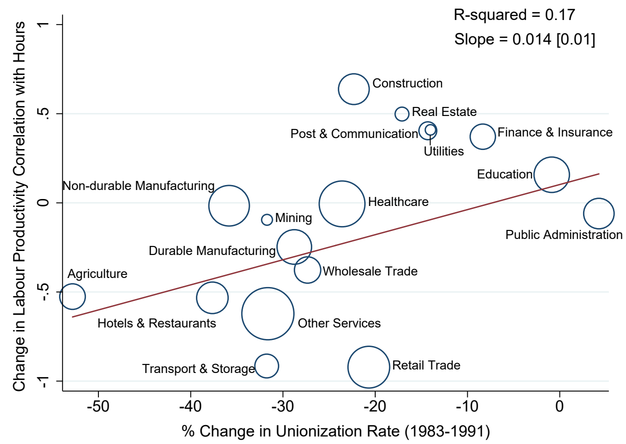

Figure 2.11: Cross-Industry Evidence for De-unionization: Productivity Correlation

Note: Data on industry-level unionization rates comes from the *Current Population Survey* (CPS), collected by Hirsch and Macpherson (2003). Labour productivity is defined as real value added per hour worked. Annual industry-level data on value-added, hours and employment between 1947 and 2010 comes from KLEMS dataset, collected by Jorgenson, Ho, and Samuels (2012). CPS industry codes for unionization and SIC industry codes in KLEMS were matched to create a consistent set of 17 U.S. industries. The Baxter and King (1999) bandpass filter between 2 and 8 years have been used to de-trend the variables. Since industry-level union data is available only from 1983 onwards, and the CPS industry codes change from 1992, to minimize concordance error I have used the change between 1983 and 1991 as the measure of change in union density. Size of the bubbles represent average industry employment level. The p-value of the slope coefficient using robust standard error is reported in parentheses.

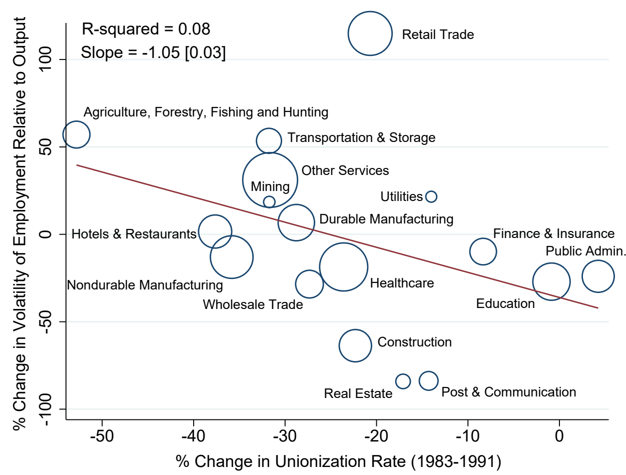

Figure 2.12: Cross-Industry Evidence for De-unionization: Relative Volatility of Employment

Note: Data on industry-level unionization rates comes from the *Current Population Survey* (CPS), collected by Hirsch and Macpherson (2003). Labour productivity is defined as real value added per hour worked. Annual industry-level data on value-added, hours and employment between 1947 and 2010 comes from KLEMS dataset, collected by Jorgenson, Ho, and Samuels (2012). CPS industry codes for unionization and SIC industry codes in KLEMS were matched to create a consistent set of 17 U.S. industries. The Baxter and King (1999) bandpass filter between 2 and 8 years have been used to de-trend the variables. Since industry-level union data is available only from 1983 onwards, and the CPS industry codes change from 1992, to minimize concordance error I have used the change between 1983 and 1991 as the measure of change in union density. Size of the bubbles represent average industry employment level. The p-value of the slope coefficient using robust standard error is reported in parentheses.

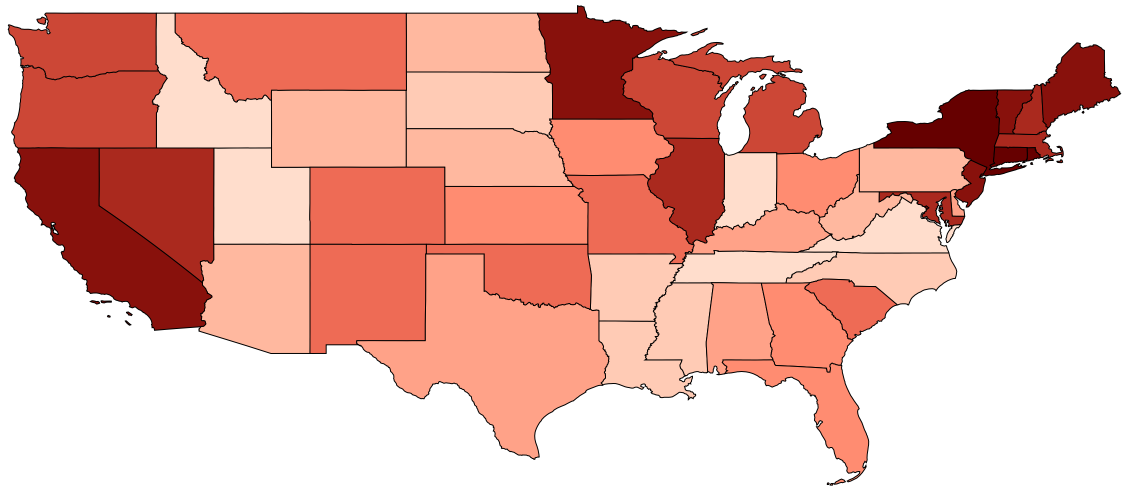

Figure 2.13: De-unionization in U.S. States around 1984
Note: 49 U.S. states (Alaska and Hawaii are missing) are grouped into deciles, with lighter shades corresponding to larger percentage of de-unionization around 1984.

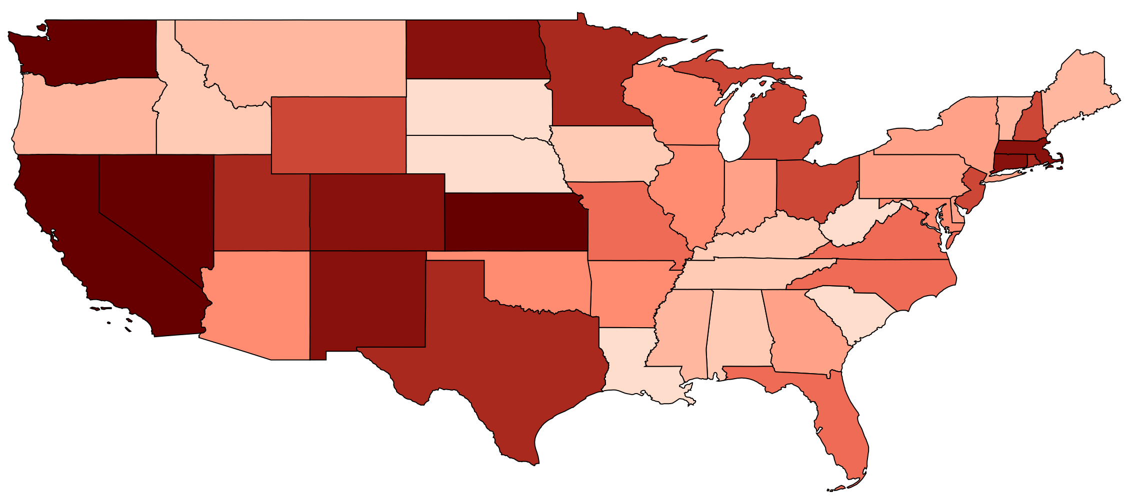

Figure 2.14: Vanishing Procyclicality of Labour Productivity in U.S. States around 1984
Note: 49 U.S. states (Alaska and Hawaii are missing) are grouped into deciles, with lighter shades corresponding to larger decrease in correlation between employment growth and output per worker growth in the pre- and post-1984 periods.

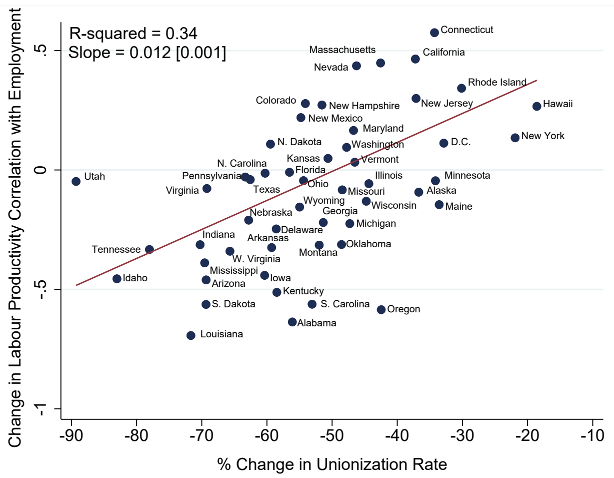

Figure 2.15: Cross-State Evidence for De-unionization

Note: Data on state-level unionization rates comes from the *Current Population Survey* (CPS), collected by Hirsch and Macpherson (2003). State-level data on real non-farm gross domestic product and total employment between 1969 and 2010 is sourced from the *Bureau of Economic Analysis* (BEA). Since hours worked data is not available at the state level, employment is used as the measure of labour input and labour productivity is defined as the state real non-farm gross domestic product per worker. I use annual growth rate as the filter because the preferred Baxter and King (1999) filter leads to 12 years of missing observations and leaves only 3 years of data before 1984. All changes in variables are calculated as the difference between the pre and post-1984 averages. Although observation for each state is weighted by its average employment in the regression, to improve readability I have not shown the weights here through bubbles, rather made it explicit in Figures 2.13 and 2.14. The p-value of the slope coefficient using robust standard error is reported in parentheses.

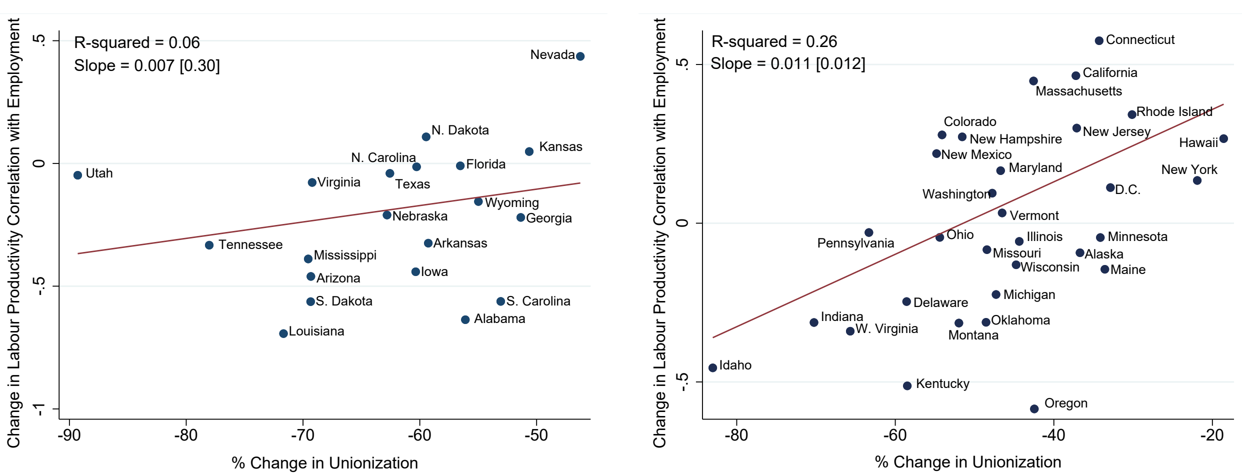

(a) Right-to-Work States

(b) Non Right-to-Work States

Figure 2.16: Cross-State Evidence for De-unionization by Right-to-Work Status

Note: Categorization of states into *Right-to-Work* and *Non-Right-to-Work* has been done based on the status in 1984. Other definitions are as in Figure 2.15.

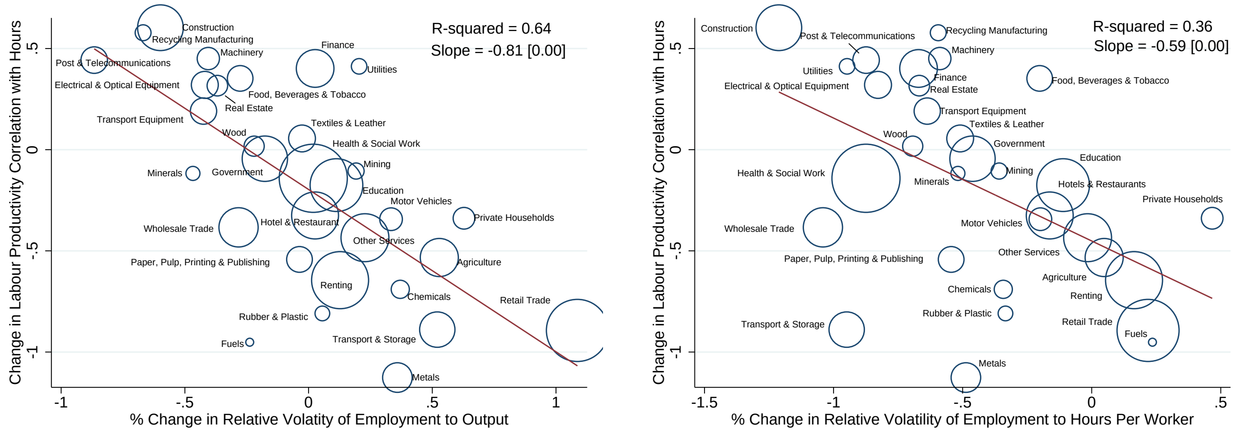

(a) Relative to Output Volatility

(b) Relative to Hours Per Worker Volatility

Figure 2.17: Relative Volatility of Employment & Change in Labour Productivity Correlation

Note: Data is for 31 U.S. industries from the annual KLEMS dataset, collected by Jorgenson, Ho, and Samuels (2012). The change in the second moments is the difference between their values in the post-1984 period (1984-2008) and the pre-1984 period (1959-1983). Regressions are weighted by the time-average of industry-level employment, depicted by the size of the bubbles. The p-value of the estimated slope using robust standard error is reported in parentheses. The Baxter and King (1999) bandpass filter between 2 and 8 years has been used to extract the cyclical component of the variables. Findings are robust to using other filters.

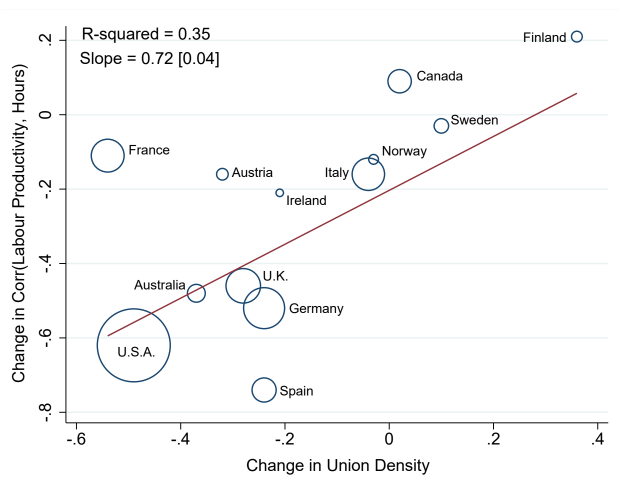

Figure 2.18: International Evidence for De-unionization

Note: All changes are between the pre- and post-1984 periods. Labour productivity is defined as real GDP per hour worked. De-trending of variables has been done using the HP-filter. Quarterly data on output and hours between 1960 and 2010 for all countries (except Spain) are taken from *OECD Economic Outlook Database*, collected by Ohanian and Raffo (2012). Annual data for Spain between 1950 and 2017 is sourced from the *Conference Board Total Economy Database*. Union density data are sourced from *OECD Annual Trade Union Density Dataset*. The regression is weighted by the time-average of country-level employment, depicted by the size of the bubbles. The p-value of the estimated slope using robust standard error is reported in parentheses.

Chapter 3

The Productivity Puzzle: Evaluating Alternative Explanations

3.1 Introduction

What explains the sudden vanishing of the procyclicality of productivity in the mid-1980s in the U.S., or the so-called *productivity puzzle*? Chapter 2 provided an answer to this question — it is the increased labour market flexibility due to rapid de-unionization. However, along with the de-regulation of the economy under President Ronald Reagan, the mid-1980s was also a time of other significant changes in the U.S. For example, the Federal Reserve under Paul Volcker made monetary policy much more accommodative (not lower interest rates much in response to negative output gaps) to curb inflation. There was a substantial decline in the importance of supply shocks for the economy as the global crude oil price came down from its peak of about $104 in April 1980 to about $22 per barrel in 1986 following conservation and insulation efforts after the Iranian Revolution. The stabilization of monetary policy, as well as the favourable oil prices and other structural factors, led to a substantial drop in the aggregate volatility of the U.S. economy from the early 1980s — the so-called Great Moderation. Around the same time, the introduction of IT capital in the workplace led to skill-biased technological change that had substantial impacts on productivity. Simultaneously, intangible capital also started becoming more important in the production process. While the service sector's share in total value-added was rising even before the 1980s, it was around the mid-1980s that the use of services as intermediate inputs in other industries picked up. The current chapter will investigate whether these changes in the shocks hitting the U.S. economy and the structural changes in economic policy and production technology have a key role to play in explaining the productivity puzzle.

The potential alternative explanations for the productivity puzzle will be assessed in two categories: first, a subset of these explanations will be studied through the lens of a quantitative model to ascertain the relative importance of these channels; and second, through separate empirical investigation outside the model framework.[33] The theoretical model will make explicit the relative quantitative importance of more accommodative monetary policy, reduced shock volatility in Great Moderation, lower hiring cost and higher bargaining power for firms due to de-unionization, and the increased importance of technology shocks relative to demand shocks. The explanatory power

[33] This has been done to keep the model simple and tractable without incorporating all the potential channels that did not have sufficient empirical evidence.

of other structural changes like the rise of the service sector, the increased relative importance of sector-specific shocks as opposed to aggregate economy-wide shocks, the increased use of intangible and IT capital, and the selective firing of least productive workers during recessions will be studied separately using a variety of data sources and empirical techniques.

I use a New Keynesian model, with only two shocks (namely, a technology shock to TFP, and a demand shock to monetary policy), that incorporates endogenous movements in labour effort or utilization with costly hiring of workers by firms. The reason to rely on the standard New Keynesian framework is that it features nominal rigidities that help generate the negative response of labour input to a positive technology shock.[34] This is a key feature of the empirical impulse response of labour input to technology shock (shown in Chapter 2) that is not captured by the RBC models that are often used in the literature to study the productivity puzzle (e.g., see Galí and van Rens (2017) and Berger (2016)). It is important to at least qualitatively get the correct sign of this impulse response because it informs us if the countercyclical response of hours was potentially driven merely by stronger technology shocks (as argued in Barnichon (2010)). Similarly, the empirically observed positive response of productivity to an expansionary demand shock is replicated in the model.

The last chapter identified de-unionization as a potential reason for the increased hiring and firing activities in the economy. This is not to claim that de-unionization is the only explanatory factor for falling employment adjustment cost. Acknowledging the possibility of multiple underlying reasons behind increased labour market flexibility, the theoretical model in this chapter will be agnostic about the exact source of the fall in labour adjustment cost along the extensive margin.[35] The process of de-unionization will be captured through two parameter-changes in the model — one, the drop in the steady-state share hiring cost in total output, and a rise in the wage bargaining power of firms. Changing each of these parameters one at a time, the model will be able to inform which channel of the impact of de-unionization is relevant in explaining the productivity puzzle. Reasonable calibration of the model can generate almost the entire drop in productivity correlations and more than 80% of the rise in relative volatility of employment. Almost all of the changes in the correlations of productivity conditional on technology and demand shocks are also qualitatively matched in the model simulations.

The rest of the chapter is organized as follows. Section 3.2 proposes a dynamic stochastic general equilibrium (DSGE) model featuring key elements of the empirical findings in Chapter 2 as discussed above. Section 3.3 then provides calibration of the model parameters, some of which are allowed to change between the pre- and post-1984 periods. Section 3.4 quantifies the performance of the model in matching the changes in business cycle moments observed in the data. Different

[34] This negative response of total hours worked or total employment to a positive technology shock can alternatively be generated under the RBC framework, provided one allows for strong habit formation in household consumption. The reason to choose the New Keynesian framework is that it also allows one to study the impact of monetary policy change, whereas in an RBC economy there is no role of money.

[35] In contrast, Zanetti (2007) shows that a standard New Keynesian monetary model with unionized labour market can explain European business cycle data much better than one with a competitive labour market, while Soskice and Iversen (2000) and Alvarez and Shimer (2014) study in details the theoretical implications of a unionized labour market with minimum wages and job-rationing.

counterfactual scenarios are also discussed here. Section 3.5 then discusses the lack of empirical evidence for a host of structural changes that could have potentially explained the productivity puzzle. Finally, Section 3.6 summarizes the key conclusions of the chapter.

3.2 Model

I will consider a New Keynesian model with two exogenous shocks — a technology shock to firm productivity, and a monetary policy shock to the nominal interest rate. I will deviate from the textbook model in two directions — first, I will explicitly consider both extensive and intensive margins of labour input adjustment (namely, employment and effort), and second, I will consider the presence of a convex cost of employment adjustment for firms. Crucially, the absence of adjustment cost along the intensive margin of effort variation will lead firms to depend more on effort adjustment when hiring costs are high. This drives the main result of vanishing procyclicality of effort, and consequently of labour productivity, in the post-1984 era when hiring costs decreased significantly.

3.2.1 Households

I assume a large number of infinitely lived identical households in the economy, with each household having a continuum of identical members represented by the unit interval. The household is the relevant decision unit for consumption and labour supply choices, and full consumption risk sharing is assumed within each household. Households seek to maximize the present value of lifetime expected utility, discounted at rate $\beta \in (0, 1)$,

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln C_t - \chi L_t\right]$$

subject to the per-period budget constraint,

$$\int_0^1 P_{it} C_{it} di + Q_t D_t \leq \int_0^1 W_{jt} N_{jt} dj + D_{t-1} + \Pi_t.$$

Here, P_{it} and C_{it} are the price and consumption of final good i, W_{jt} is the nominal wage paid at firm j, D_t denotes the amount of one-period bonds purchased at price Q_t, and Π_t represents any lump-sum income including dividends from ownership of firms and government taxes and transfers. Household's aggregate consumption bundle, $C_t \equiv \left(\int_0^1 C_{it}^{\frac{\varepsilon-1}{\varepsilon}} di\right)^{\frac{\varepsilon}{\varepsilon-1}}$ is an index of the quantities consumed of different types i of final goods, and is priced at $P_t \equiv \left(\int_0^1 P_{it}^{1-\varepsilon} di\right)^{\frac{1}{1-\varepsilon}}$, with $\varepsilon > 1$ being the Kimball aggregation parameter for the unit mass of final goods. The second term in the period utility function represents disutility from effective labour supply L_t, which not only depends on the fraction N_t of household members who are employed but also the amount of effort, E_t exerted by each employed member. More specifically, I assume the following functional form for effective labour supply, $L_t \equiv \left(\frac{1+\zeta E_t^{1+\phi}}{1+\zeta}\right) N_t$. The parameter $\chi > 0$ measures the importance of disutility from forgone leisure, while $\zeta \geq 0$ measures the importance of effort in that disutility from working. The elasticity parameter $\phi \geq 0$ measures the degree of increasing marginal disutility from exerting more effort.

I make the simplifying assumption of constant hours per worker so that the only source of intensive margin adjustment in labour supply is effort. More importantly, I assume that households take into account the endogenous impact of employment adjustment decisions on the level of effort exerted by each of its members.

Consumption maximization for any given level of expenditure, $P_t C_t$ is done by choosing the optimal amount of consumption of each intermediate good, and the resulting demand function for good $i \in [0, 1]$ is given by

$$C_{it} = \left(\frac{P_{it}}{P_t}\right)^{-\varepsilon} C_t \tag{3.1}$$

The intertemporal optimality condition is given by

$$Q_t = \mathbb{E}_t \left(\frac{P_t}{P_{t+1}} \Lambda_{t,t+1}\right) \tag{3.2}$$

where $\Lambda_{t,t+k} \equiv \beta^k \frac{C_t}{C_{t+k}}$ $\forall t, k$ is the stochastic discount factor measuring the marginal rate of intertemporal substitution.

3.2.2 Firms

I model the production side of the economy as a two-sector structure — final and intermediate goods sectors. Households supply labour only to firms in the intermediate goods sector who produce a variety of intermediate goods. Final goods firms do not employ labour, and effectively only repackage the intermediate goods and sell them in the market at a mark-up over marginal cost, subject to restrictions in the frequency of their price-setting decisions.

Final Goods

A continuum of monopolistically competitive firms constitutes the final goods market, with each firm $i \in [0, 1]$ producing a differentiated final good Y_{it} according to the production function, $Y_{it} = X_{it}$, where X_{it} is the quantity of the single intermediate good used by the final good firm i as an input. In the absence of nominal rigidities, profit maximization leads to the following price-setting condition for all t,

$$P_{it} = \left(\frac{\varepsilon}{\varepsilon - 1}\right) P_t^I \tag{3.3}$$

where P_t^I is the price of the intermediate good, and the factor $\left(\frac{\varepsilon}{\varepsilon-1}\right)$ is the optimal mark-up over the marginal cost of production. However, à la Calvo (1983), I assume that final goods firms are precluded from setting their prices optimally in any period with probability $\theta_p \in [0, 1]$. This probability is independent both across firms, and of the time elapsed since the last nominal adjustment. This ensures that the fraction of firms changing their prices in any given period is a constant $(1 - \theta_p)$, which can be interpreted as the degree of nominal flexibility in the economy.

Thus, the law of motion for the aggregate price level in the economy, P_t becomes a weighted average of the optimally chosen price, P_t^* and the price that prevailed in the last period, P_{t-1}, with the weight being the probability of nominal adjustment:

$$p_t = \theta_p p_{t-1} + (1 - \theta_p)\, p_t^* \tag{3.4}$$

where the lower case letters denote the natural logarithms of the corresponding upper case variables. Since all firms face an identical problem every period, the optimal price, P_t^* is the same across firms, and is given by

$$p_t^* = \mu^p + (1 - \beta\theta_p) \sum_{k=0}^{\infty} (\beta\theta_p)^k \, \mathbb{E}_t \left(p_{t+k}^I\right) \tag{3.5}$$

where $\mu^p \equiv \ln\left(\frac{\epsilon}{\epsilon - 1}\right)$. Combining equations (3.4) and (3.5), one can derive the inflation equation as follows:

$$\pi_t^p = \beta \mathbb{E}_t \left(\pi_{t+1}^p\right) - \lambda_p \hat{\mu}_t^p \tag{3.6}$$

where $\pi_t^p \equiv p_t - p_{t-1}$ is price inflation, $\lambda_p \equiv \frac{(1-\theta_p)(1-\beta\theta_p)}{\theta_p}$ and $\hat{\mu}_t^p \equiv \mu_t^p - \mu^p = p_t - p_t^I - \mu^p$ is the deviation in logs of the average mark-up from its steady state value.[36]

Intermediate Goods

Each perfectly competitive intermediate goods firm $j \in [0, 1]$ faces the production function $Y_{jt}^I = A_t \left(E_{jt}^{\psi} N_{jt}\right)^{(1-\alpha)}$, where A_t is the technology term common across all firms, the parameter $\psi \in (0, 1)$ measures the additional returns to effort over employment, and $\alpha \in (0, 1)$ denotes non-labour income share in the economy.[37] In the calibration in Section 3.3, the parameters ψ and α are chosen such that they satisfy non-increasing returns to scale of the production function: $(1 - \alpha)(1 + \psi) \leq 1$. The productivity term A_t has the following exogenous stochastic process: $a_t \equiv \ln(A_t) = \rho_a a_{t-1} + \varepsilon_t^a$, where ε_t^a is a white noise process with variance $\sigma_{\varepsilon^a}^2 > 0$. Since the production function explicitly includes the factor utilization term, namely, effort E_{jt}, the productivity term A_t should be interpreted as the utilization-adjusted TFP.

Labour Market

Workers get separated from their jobs at intermediate goods firms at the exogenous gross rate of $\delta \in (0, 1)$, but every period t firm j hires back new workers H_{jt} $\forall j$, subject to a per-worker

[36] See Galí (2008) for derivation of equations (3.5) and (3.6).

[37] One might be concerned that whatever is being labelled as 'effort' in the production function is in fact capital, the missing factor of production. In Appendix B.2, I contrast the cyclical properties of capital with that of factor utilization (which is a proxy for 'effort') and show how they evolved differently. This allays the identification concern of 'effort' being equivalent to capital. Empirically, it has so far been impossible to distinguish between capital utilization and worker utilization rates, e.g., Fernald (2014) uses hours per worker as the proxy for both labour and capital utilization, and the capacity utilization measure by the Federal Reserve is a combined measure of the intensive margin of all factors of production. Given this lack of identification of the intensive and extensive margins of labour and capital separately, I do not include capital in the analysis because it would not be possible to separately identify time-variation in capital adjustment cost and employment adjustment cost.

adjustment cost function, $G_t = \Gamma H_t^{\gamma}$, with the parameters Γ and γ being strictly positive, and $H_t \equiv \int_0^1 H_{jt} dj$ denoting aggregate level of hiring. The assumption of a hiring cost as opposed to a firing cost is dictated by the simplicity of calibration. There is no reason to believe that the hiring and firing costs are symmetric, and this assumption should be relaxed in future research. Nevertheless, the motivation for this hiring cost is best summarized in Heckman, Pagés-Serra, Edwards, and Guidotti (2000): *"...in the face of a positive shock firms may want to hire additional workers, but they will take into account that some workers may have to be fired in the future if demand turns down. This prospective cost acts as a hiring cost..."* Moreover, more powerful unions can make this hiring cost to rise for firms. This link between union density and hiring cost will be crucial later for the calibration of the model.[38] However, there is no a priori reason to believe that de-unionization could be the only reason for a decline in the hiring cost. As discussed in Chapter 2, increased use of temporary workers or the advent of online job-search platforms can also be thought of as potential determinants of hiring cost. Nonetheless, the presence of the job separation rate and the hiring by the firm implies that employment at firm j has the following law of motion

$$N_{jt} = (1 - \delta) N_{jt-1} + H_{jt} \tag{3.7}$$

Because of the presence of labour market frictions in the form of a hiring cost, wages and employment may differ across firms, since they cannot be instantaneously arbitraged out by the free movement of workers from low to high wage firms. Therefore, in what follows, the subscript j on wages and employment will signify this potential difference across firms. Faced with the common hiring cost function, G_t and given the nominal wage W_{jt}, firm j's optimal hiring policy is given by the condition,

$$MRPN_{jt} = \frac{W_{jt}}{P_t} + G_t - (1-\delta)\mathbb{E}_t \left(\Lambda_{t,t+1} G_{t+1}\right) \tag{3.8}$$

where $MRPN_{jt} = (1-\alpha)(1-\Psi_F) \frac{P_t^I}{P_t} \frac{Y_{jt}^I}{N_{jt}}$ is the marginal revenue product of employment expressed in terms of final goods. The non-zero term $\Psi_F \equiv \frac{\alpha\psi}{1+\phi-(1-\alpha)\psi}$ arises due to the endogenous response of effort to changes in employment. This condition implies that each period the firm hires workers up to the point where the marginal revenue from an additional employment equals the cost of that marginal worker, where the cost involves not only the wage and the hiring cost in the current period, but also the discounted future savings from having to hire $(1-\delta)$ fewer workers in the following period. Solving equation (3.8) forward, one has the following expression for the average hiring cost,

$$G_t = \mathbb{E}_t \left[\sum_{k=0}^{\infty} \Lambda_{t,t+k} (1-\delta)^k \left(MRPN_{jt+k} - \frac{W_{jt}}{P_{t+k}} \right) \right] \tag{3.9}$$

[38] Freeman and Medoff (1982) highlight the substitution away from production workers towards other factors of production in the presence of higher labour costs in unionized manufacturing.

For notational convenience in deriving the log-linearized version of equation (3.8) later on, I define the *net* hiring cost as $B_t \equiv G_t - (1-\delta)\,\mathbb{E}_t\left(\Lambda_{t,t+1} G_{t+1}\right)$, such that equation (3.8) can be re-written as

$$MRPN_{jt} = \frac{W_{jt}}{P_t} + B_t \tag{3.10}$$

I assume wages are negotiated and potentially adjusted every period through a Nash bargaining process between the intermediate goods firms and the households to split the total surplus generated from an established employment relation. The surplus accruing to the firm j and the household members who work at firm j are given by the following two equations respectively,

$$S^F_{jt} = MRPN_{jt} - \frac{W_{jt}}{P_t} + (1-\delta)\,\mathbb{E}_t\left(\Lambda_{t,t+1} S^F_{jt+1}\right) \tag{3.11}$$

$$S^H_{jt} = \frac{W_{jt}}{P_t} - MRS_{jt} + (1-\delta)\,\mathbb{E}_t\left(\Lambda_{t,t+1} S^H_{jt+1}\right) \tag{3.12}$$

where $MRS_{jt} = \frac{\chi C_t}{1+\zeta} + \Psi_H \frac{P^I_t}{P_t}\frac{Y^I_{jt}}{N_{jt}}$ is the household's marginal rate of substitution between consumption and employment at firm j, or equivalently the marginal disutility of employment expressed in terms of the final goods bundle. The non-zero term $\Psi_H \equiv \frac{\psi}{1+\phi}\left(1 - \frac{(1+\phi)W_{jt}N_{jt}}{(1+\phi-\psi)P_t C_t}\right)$ arises due to the endogenous response of effort to changes in employment. It is interesting to note that profit maximization by firms implies that the firm surplus S^F_{jt} equals the per worker hiring cost, G_t. The average hiring cost can thus be interpreted as what the firm potentially saves from maintaining an existing employment relation.

Denoting the relative bargaining power of firms vis-à-vis workers by the parameter $\xi \in (0,1)$, the Nash bargaining set-up solves the following problem

$$\max_{\{W_{jt}\}} \left(S^F_{jt}\right)^{\xi}\left(S^H_{jt}\right)^{1-\xi}$$

subject to equations (3.11) and (3.12). The solution to the above bargaining problem implies a constant share rule, $\xi S^H_{jt} = (1-\xi)\,S^F_{jt}$, which translates to the equilibrium wage condition, $\frac{W_{jt}}{P_t} = \xi MRS_{jt} + (1-\xi)\,MRPN_{jt}$. Substituting for $MRPN_{jt}$ and MRS_{jt} in the above wage equation, one can write the average Nash-bargained wage (upto a first order approximation) as

$$\frac{W_t}{P_t} = \xi MRS_t + (1-\xi)\,MRPN_t \tag{3.13}$$

So far, I have assumed that firms can re-negotiate wages every period. However, this militates against the empirical evidence of substantial nominal wage rigidities, e.g., the average frequency of wage changes is often found to be more than one year. Incorporating Calvo-type nominal wage stickiness is quite straight-forward and does not alter the main intuition of the wage-setting process discussed here. In particular, I assume θ_w fraction of firms cannot re-optimize their nominal wages in each period, thereby leading to the law of motion for average nominal wage, $w_t \equiv \int_0^1 w_{jt}dj$ as $w_t = \theta_w w_{t-1} + (1-\theta_w)\,w^*_t$. In this set-up, the deviation between the actual average real wage ($\omega_t \equiv w_t - p_t$) and the average Nash-bargained wage under a counterfactual flexible wage environment $\left(\omega_t^{target}\right)$ drives the wage inflation in the economy.[39]

3.2.3 Monetary Policy

I assume a standard Taylor-type interest rate rule for the Central Bank,

$$i_t = \rho i_{t-1} + (1-\rho)\left(\phi_\pi \pi_t^p + \phi_y \hat{y}_t\right) + \phi_{\Delta y}\Delta\hat{y}_t + \nu_t \tag{3.14}$$

where $i_t \equiv -\ln Q_t$ is the nominal yield on a one-period riskless bond, ρ is the persistence in monetary policy, $\hat{y}_t$ is the logarithm of the period t output gap in the economy, and ν_t is the exogenous policy shifter. The monetary policy shock ν_t is assumed to follow an AR(1) process: $\nu_t = \rho_\nu \nu_{t-1} + \varepsilon_t^\nu$, where the persistence parameter $|\rho_\nu| < 1$ and ε_t^ν is a white noise process with variance $\sigma_\nu^2 > 0$. The degree to which the Central Bank accommodates exogenous shifts in productivity partly determines the coefficient of the output gap in the Taylor rule. In particular, smaller the parameter ϕ_y, the more accommodating is the monetary policy. Since I have already shown empirically that the response of hours and employment turned less countercyclical or sometimes even procyclical after 1984, one can expect to see the parameter ϕ_y turning smaller in magnitude in the later years. It should be noted that a countercyclical response of employment to a technology shock is contingent on the monetary policy being not too accommodative.

3.2.4 Equilibrium Conditions

I assume that hiring costs take the form of a bundle of final goods given by the same aggregation as the one defining the consumption index. This implies that the demand for each final good is given by $Y_{it} = \left(\frac{P_{it}}{P_t}\right)^{-\varepsilon}(C_t + G_t H_t)$. The goods market clearing condition is thus given by

$$Y_t \equiv \left(\int_0^1 Y_{it}^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}} = C_t + G_t H_t \tag{3.15}$$

The aggregate relation between final goods and intermediate input is given by

$$X_t \equiv \int_0^1 X_{it} dj = D_t^p Y_t \tag{3.16}$$

where the chasm $D_t^p \equiv \int_0^1 \left(\frac{P_{it}}{P_t}\right)^{-\varepsilon} di \geq 1$ between the quantities produced and consumed of the different final goods arise due to the price dispersion caused by nominal rigidities. However, in the neighbourhood of the zero-inflation steady state, $D_t^p \simeq 1$, and hence the aggregate production function can be approximated by the following condition,

$$Y_t = A_t \left(E_t^\psi N_t\right)^{(1-\alpha)} \tag{3.17}$$

[39] For the system of log-linearized equations dictating the wage-setting process with nominal rigidity, refer to equations (B.11) through (B.13) in Appendix B.1.

3.3 Calibration

Having put in place a DSGE framework with endogenous effort choice and costly employment adjustment, I now study the quantitative performance of that model.[40] Specifically, I will calibrate the parameters of the model to reasonable values often estimated or assumed in the literature, and then check whether structural changes in some of them between the pre and post-1984 periods can generate the empirically observed changes in the business cycle moments.

A typical DSGE model, like the one presented above, contains a lot of parameters. For the ease of exposition, I discuss the calibration of the entire set of parameters in four groups: (i) parameters affected by de-unionization, namely, the share of hiring cost in GDP, Θ and the wage bargaining power, ξ; (ii) the accommodative stance of monetary policy, ϕ_y, which changed during the Volcker-era and had an impact on the economy's response to technology shocks; (iii) parameters pertaining the volatility of the exogenous shocks to technology and monetary policy, namely, σ_a and σ_ν, which decreased during Great Moderation; and (iv) other parameters that I will consider to have remained stationary over the period under study.

3.3.1 Structural Changes due to De-unionization

A fall in hiring cost leading to a rise in the relative dependence on the extensive margin of labour adjustment is the key mechanism under study here. Denoting by Θ the steady-state share of total hiring cost in real output, i.e., $\Theta \equiv \frac{\bar{G}\bar{H}}{\bar{Y}}$, the main hypothesis can thus be captured by a decrease in Θ in the post-1984 period. I consider a fall in the share of the hiring cost in GDP from 3% in the pre-1983 period to 1% in the post-1984 era. These magnitudes are in line with the estimates of the hiring cost share by Silva and Toledo (2009) and used for calibration in Hagedorn and Manovskii (2008). They estimate hiring cost to be roughly 4.5% of the average quarterly wage. Assuming average wage to be 67% of real output (which is nothing but the labour share in total compensation), the hiring cost as a share of GDP is calibrated to be 3% in the pre-1984 period. Now, union membership rate in private non-farm U.S. industries was about 21% in 1979 after which time it started falling sharply, and reached 1/3 of that value at roughly 7% by 2009. I, therefore, calibrate the hiring cost share in GDP in the post-1984 era as 1/3 of its pre-1984 value of 3%.[41]

De-unionization not only affects the hiring cost of workers but also increases the bargaining power of firms. Assuming equal bargaining power, i.e., $\xi = 0.50$, for calibration purposes, is the standard in the literature. Felix and Hines Jr. (2009) find that workers in a fully unionized firm capture roughly 54% of the benefits of lower state corporate income tax rates in the U.S., which roughly indicates an equal bargaining power between workers and firms. Starting from an equal

[40] For ease of exposition, I have collected in Appendix B.1 all the model equations with the variables being measured in logarithms of deviations from their zero-inflation steady-state values.

[41] Galí and van Rens (2017) also consider a fall in the share of hiring cost from 3% to 1% of GDP, but their calibration choice is motivated by the fall in the gross job separation rate. However, as argued in Appendix 2.4.2, the reduction in job separation rate appears to be an unlikely explanation for the fall in hiring costs when international evidence is taken into account. Moreover, data on quarterly job flows from Shimer (2012) show that job separation rate fell by only about 10% in the post-1984 era.

bargaining power between workers and firms in the pre-1983 period, I allow the parameter to increase by 67% in the post-1984 period to 0.84, mirroring the fall in union density in the private nonfarm business sector in the U.S.

3.3.2 Monetary Policy Change

To capture the reduced response of hours on the impact of a technology shock in the post-1984 period, it is crucial to allow for parameter ϕ_y, capturing the degree of accommodation of technology shocks by the monetary policy, to fall in the post-1984 era. I use the values in Smets and Wouters (2007) who estimate the Taylor-rule parameters separately for two periods: 1966 through 1979, and 1984 through 2004, and find that ϕ_y decreased from 0.17 to 0.08 between the two periods.[42] As mentioned earlier, a more accommodative monetary policy counteracts the fall in productivity correlation with hours. In other words, the fall in productivity correlations would have been even larger had there been no structural change in the monetary policy.

3.3.3 Exogenous Shocks: Changes during Great Moderation

For comparability to the literature, I will calibrate the persistences of the productivity and monetary policy shocks to values typically assumed in the literature. In particular, the persistence of the technology shock is assumed to be 0.90, following Galí (2011), and that of the monetary shock is fixed at 0.50, following Galí (2011) and Barnichon (2010). Accounting for the fall in volatility due to Great Moderation, I calibrate the post-1984 standard deviation of the technology shock to be 70% of the pre-1983 value, and that of the monetary shock to be 50% of its corresponding pre-period value. This choice of reductions of 30% and 50% in the standard deviations of the technology and demand shocks is motivated by the findings in Barnichon (2010).[43] Allowing for the shock variances to change in a two-variable SVAR with labour productivity growth and unemployment, Barnichon (2010) finds these magnitudes of drops in the volatilities of technology and demand shocks around the mid-1980s. The calibrated values for the volatilities in the two periods are reported in Table 3.1.

3.3.4 Stationary Parameters

The fourth set of parameters correspond to those which are quite standard in the literature and are arguably not likely to have changed significantly between the pre and post-1984 periods. The complete list of these parameters and their calibrated values are presented in Table 3.2. While most of the parameters are calibrated to some well-established estimates in the literature, the last three parameters in Table 3.2 are somewhat arbitrarily chosen because of lack of consensus in the

[42] It should be noted that the positive response of productivity to an expansionary demand shock (i.e., a negative Taylor-rule shock to the interest rate) is contingent on the monetary policy being not too accommodative. This condition is maintained by the set of estimates in Smets and Wouters (2007).

[43] As external evidence, in Appendix B.3 I show that the volatility of the monetary shock as measured by Romer and Romer (2004) empirically decreased by about 50%.

empirical literature — the curvature of the convex average hiring cost function (γ), the degree of increasing marginal disutility from higher effort (ϕ), and the additional curvature of effort in the production function (ψ). However, as a robustness check in Section 3.4.3, I show that assuming different values for these three parameters do not significantly alter the main quantitative findings from the model.

The calibration of the production function parameters ψ and α needs attention. Suppose $\alpha = 0$. Then the production function $Y_{jt} = A_t E_{jt}^{\psi} N_{jt}$ can be thought as a special case of a standard Cobb-Douglas production function with effective employment, $E.N$ and capital K, $\tilde{Y}_{jt} \equiv A_t \left(E_{jt} N_{jt}\right)^{\psi} K_{jt}^{(1-\psi)}$, provided capital per worker, $\frac{K_{jt}}{N_{jt}}$ is held constant over the business cycle, since $\tilde{Y}_{jt} = Y_{jt} \left(\frac{K_{jt}}{N_{jt}}\right)^{1-\psi}$. This parameterization of the production function assumes short-run increasing returns to effective labour, which is a standard feature of models trying to generate procyclical movements of labour productivity in response to a demand shock.[44] In this Cobb-Douglas representation, the parameter ψ has the natural interpretation of the labour share in aggregate production, and can, therefore, be calibrated to a value of 0.67. This route of calibration has been adopted in Barnichon (2010). However, the increasing returns to scale when $\alpha = 0$ is inconsistent with the perfectly competitive nature of the intermediate goods firms considered here. As an alternative, one can interpret α as the share of non-labour inputs in the total factor cost of production and can be calibrated to roughly 0.33.[45] This leaves the calibration of ψ open within certain limits — ψ cannot exceed 0.50 to avoid increasing returns to scale.

3.4 Quantitative Performance of Model

Having calibrated the model parameters separately for the pre and post-1984 sub-periods, I will now examine how well the model can match the main phenomenon of vanishing procyclicality of labour productivity, along with other changes in business cycle moments like the rising relative volatility of employment, the falling procyclicality of real wages, etc.

3.4.1 Business Cycle Moments

There are multiple parameter changes in the calibration for the two sub-periods. It is therefore natural to ask what role does each parameter change play in explaining the differences in the business cycle moments. For that purpose, I will introduce the parameter changes in Table 3.1 one at a time.

[44] Gordon (1993) emphasizes the theoretical need for the presence of such increasing returns to labour for explaining business cycle facts. Empirical works by Basu (1996), Basu, Fernald, and Kimball (2001) and others have confirmed increasing returns to scale production functions, both for durable manufacturing and services industries.

[45] The widely known empirically observed drop in the labour share of output or total compensation in the U.S. accelerated only from the early 2000s and therefore is not considered as a candidate for explaining the productivity puzzle of the mid-1980s.

De-unionization. In Table 3.3, I first see how de-unionization alone performs in capturing the changes in the moments. Column (1) reports the empirically observed changes in business cycle moments between pre and post-1984 periods, while column (4) reports the total change explained by de-unionization. However, since a fall in union density is captured by two parameter-changes, namely, a fall in the share of hiring cost in GDP, Θ and a rise in the firms' bargaining power, ξ, I also show the relative contribution of these two channels in the total effect of de-unionization in columns (2) and (3) respectively. Comparing columns (1) and (4) in Table 3.3 one can see that the parameter changes attributed to de-unionization perform well in matching the empirically observed drop in productivity correlations, both for unconditional correlations as well as conditional on technology and demand shocks. For the relative volatility of employment, the baseline calibration of the model captures more than 80% of the total rise in the data. Also, most of the changes in these moments can be attributed to the change in the hiring cost parameter, which is the central mechanism discussed here.

Finally, the cyclical properties of real wages have also changed in the U.S. around the same time as the productivity puzzle and the Great Moderation. The model's ability to capture the changes in the cyclical wage correlations is primarily driven by the change in the bargaining power parameter. This importance of the bargaining parameter in determining wage dynamics is not surprising, given that the parameter directly enters the real wage equation (3.13). Regarding the volatility of real wages, the current model predicts a fall in the post-1984 era. However, empirical evidence on wage volatility has been mixed. Champagne, Kurmann, and Stewart (2017) discuss how average hourly wage volatility in the U.S. has diverged across different data sources: the Labor Productivity and Costs (LPC) program, the Current Population Survey (CPS), and the Current Employment Statistics (CES). Supplements and irregular earnings of high-income workers, included only in the LPC, drive the rising volatility in LPC earnings as opposed to CPS and CES based measures. One way to match the rising volatility of real wages (e.g., Champagne and Kurmann (2013) and Nucci and Riggi (2013)) in the current model would be to introduce real wage rigidity and let it decline in the post-1984 period.[46] This can be done by introducing wage indexation to past inflation, or though endogenous real wage rigidity that depends on the size of the wage bargaining set in equilibrium. Lemieux, Macleod, and Parent (2009) discuss the rising importance of performance wages in the U.S. economy, which can also help explain the rising wage volatility. These channels are however absent in the current version of the model and remain a task for future research.

Accommodative Monetary Policy. In Table 3.4, column (3) shows that more accommodative monetary policy by the Federal Reserve cannot induce large changes in the productivity moments,

[46] A real wage rigidity decline in the post-1984 period is not to be confused with increasing nominal wage rigidity during the same period. Using Bayesian estimation, Smets and Wouters (2007) find nominal wage rigidity to have gone up after the mid-1980s, although the increase is not statistically significant. Under increasing nominal wage rigidity, firms cannot change wages as frequently as they want and need to rely on adjusting employment more. This channel of more employment adjustment further depresses the procyclicality of productivity and increases the relative employment volatility. This mechanism is highlighted in Gu and Prasad (2018).

and most of those changes go against the empirically observed direction of moment changes. As argued in Section 2.3.2, allowing for a more accommodating monetary policy means that conditional on a positive technology shock when the output gap increases, the contraction induced through monetary policy is less severe. This implies that with a lower ϕ_y, output and employment increases more in response to a given positive technology shock, thereby increasing the cyclical correlation of productivity with labour input. This corroborates the empirical finding in Section 2.3.2 that the negative impulse response of hours worked to a positive technology shock is somewhat muted after the mid-1980s. To summarize, in absence of the more accommodative stance of the Federal Reserve under Volcker, the drop in productivity correlations would have been even more severe.

Reduction in Shock Volatility. Column (4) of Table 3.4 shows that the model's ability to match the changes in the business cycle moments is not contingent on the drop in volatilities of the exogenous shocks during Great Moderation. There are two aspects to this observation. First, a uniform reduction of volatilities of shocks per se cannot be expected to change the correlation between variables. In that sense, this finding is not surprising. However, in the calibration, the reduction in technology shock volatility was smaller than the fall in demand shock volatility. This mechanically increases the importance of technology shock in the post-1984 period. Since technology shocks induce countercyclicality of productivity with labour input, this should explain part of the vanishing procyclicality of productivity. This mechanism was highlighted in Barnichon (2010). Nevertheless, one can see from column (4) that even this channel of non-uniformity in volatility reduction could not explain any significant amount of the productivity puzzle.

One of the main indicators that substantiated the role of increased flexibility in the labour market as the key driving factor behind the labour productivity puzzle was the drastic change in the correlations of productivity *conditional* on demand shocks. If it were merely the case of changing relative importance of technology and demand shocks that explained the fall in the unconditional productivity correlations (as argued in Barnichon (2010)), then structural changes like the decline of hiring costs would not have been the significant channels for explaining the puzzle. The substantial fall in the productivity correlations conditional on the monetary policy shock in the model corroborates that empirical finding.

The finding that the fall in the productivity correlations conditional on a demand shock is driving the unconditional moments implies that demand shock should be the main source of variation for output and employment dynamics over the business cycle. This is empirically corroborated by the dominance of non-technology shocks in explaining the total cyclical volatility of per capita hours in Figure 3.2. Since the only non-technology or demand shock in the model is the monetary policy shock, it is the dominant source of business cycle variation here. However, it should be noted that Smets and Wouters (2007) find that in the presence of a variety of demand shocks, e.g., exogenous spending shock, risk premium shock, investment-specific technology shock, etc., the role of monetary policy shock is quite limited in the cyclical variation of output. Thus, the predominant role played by the monetary policy shock in this model should be thought as a consequence of the loading of all variation due to various demand shocks onto a single monetary policy shock.

3.4.2 Impulse Response Functions

One check for the quantitative validity of the model is to be able to generate impulse responses that are in line with the empirically observed ones. Figure 3.1 shows how the impulse response of employment rate to a positive technology shock (Panel (a)), and that of average labour productivity to a contractionary monetary policy shock (Panel (b)) have both become muted in the post-1984 period. These changes in model-implied impulse responses are indeed qualitatively the same as those observed empirically (discussed in Section 2.3.2).

It is interesting to know what parameter changes in the model are driving the changes in the impulse responses to shocks in the post-1984 period. While the muted negative response of employment to technology shock is almost entirely driven by the fall in ϕ_y (which implies a more accommodative monetary policy in the post-1984 era), the reduced magnitude of the rise in productivity due to a contractionary monetary policy shock is caused by the fall in hiring cost Θ. These changes in impulse responses once again prove that the labour productivity puzzle cannot be explained by the rise in the relative importance of technology shock (because the countercyclical property of technology shocks became muted post-1984), rather by structural changes in the economy that caused productivity to respond less to demand shocks over the business cycle.

3.4.3 Robustness to Parameter Calibration

The hiring cost function is taken to be quadratic in the baseline calibration of the model. However, there is no agreement in the literature as to the degree of convexity of the function. Mortensen and Nagypál (2007) find that in the presence of search frictions with linear vacancy posting costs, the matching function has an unemployment elasticity of 0.6. Interpreting employment adjustment costs as search frictions, a natural calibration for γ in the current model is therefore 0.6. On the other hand, Merz and Yashiv (2007) directly estimate the convexity of the average employment adjustment cost and reports a value of 2.4. In Table 3.5, I show the robustness of the quantitative model predictions for different values of γ in this range.

In the baseline calibration of the model, the degree of increasing marginal disutility from exerting more effort, ϕ is taken to be 1 and the additional curvature of effort in the production function, ψ is taken to be 0.5. Since there is no consensus in the literature regarding these values, in Tables 3.6 and 3.7, I show the robustness of the model's quantitative performance for alternative values of ϕ and ψ.

Finally, the nominal rigidities in the baseline calibration has been assumed to have remained constant across the pre- and post-1984 periods. However, Smets and Wouters (2007) show that there has been a significant rise in the price rigidity for goods in the post-1984 period because of the reluctance of firms to change prices in an era of low inflation under the Great Moderation. Therefore, as a robustness check, I show in Table 3.8, that allowing for the price and nominal wage rigidities to change between the two sub-periods according to the estimates in Smets and Wouters (2007) does not qualitatively alter the findings.

3.5 Other Plausible Explanations: Lack of Evidence

Having provided a coherent structural explanation for a host of changes that occurred in the correlation and volatility patterns of key economic variables around the mid-1980s in the U.S., I will now try to argue that some of the other plausible channels that have been explored in the literature as potential explanations for the productivity puzzle do not hold up to closer empirical scrutiny.

3.5.1 Vanishing Countercyclicality of Labour Quality

One simple explanation for the increased countercyclicality of productivity could be that the firms are selectively firing their least productive workers during recessions, and they are doing this more intensely after the mid-1980s because of either a greater ability to measure individual worker productivity (possibly due to the availability of better monitoring technology) or greater ease of hiring and firing workers (possibly due to factors like de-unionization). Simply put, if firms fire their least productive workers in recessions, the average productivity of the workers remaining in the workforce automatically rises in bad times, driving the countercyclicality of productivity. This channel has been highlighted in Berger (2016). To ascertain whether this is indeed the case, I compute the cyclical correlation of a measure of labour quality with business sector output. The measure of labour quality used is the Labour Quality Index constructed by Aaronson and Sullivan (2001) from 1979 onwards using CPS data on individual worker wage, sex, job experience and education, while the pre-1979 data is the annual BLS Multi-Factor Productivity estimate of labour composition interpolated by Fernald (2014) using the method outlined in Denton (1971). Plotting the rolling window correlation of labour quality and output along the business cycle in Figure 3.4, I find that while it is true that labour quality rises in recessions (as manifested by the negative cyclical correlation of labour quality with output), there is no evidence that this phenomenon has intensified in the post-1980s period (there is no discernible difference in the correlation before and after the 1980s).[47] This implies that the greater ease of hiring and firing workers did not translate into a more selective firing of low-quality workers during recessions, rather more firing in general of all workers (or the 'average-quality representative' worker).

[47] Studying individual worker productivity from a large technology-based services company in the U.S. between 2006 and 2010, Lazear, Shaw, and Stanton (2016) find that most of the increased productivity during the Great Recession was driven by an increased effort of the retained workers, and neither by the selective firing of least productive workers nor by hiring more productive new workers. Their finding corroborates the importance of the effort channel that is highlighted in this thesis, as opposed to the labour quality channel.

3.5.2 Rise of Service Sector

The rise of the service sector and the corresponding decline in the share of value-added and employment for the manufacturing industries could have led to the fall in the labour productivity correlations.

One possible channel is the so-called *composition effect* — if labour productivity in the service sector is less correlated with output and hours (arguably due to more flexible work hours), then a simple compositional shift in the share of value-added or employment towards services can explain the decline in the aggregate productivity correlations. However, the labour productivity correlations in Table 3.9 clearly show that the two sectors had strikingly similar correlations even before the mid-1980s, and both of them experienced a similar drop in labour productivity correlations over the business cycle. Moreover, this compositional shift towards services has been too gradual (see Panels (a) and (b) of Figure 3.5) to explain the sudden drop in the productivity correlations. This refutes the claim that a simple compositional shift towards a service economy was responsible for the vanishing procyclicality of aggregate productivity.

The second channel through which the rise of the service sector can contribute towards falling aggregate productivity correlations is the *substitution effect* — if there is a larger share of services intermediate inputs in the economy then the labour productivity of all sectors will mimic that of the services sector.[48] While there has indeed been an increase in the share of services intermediate inputs since the early 1980s (see Panel (c) of Figure 3.5), it is not the case that the manufacturing sector started to have similar productivity correlations as the service sector only after the increased use of services intermediate inputs. Moreover, looking at a cross-section of 31 U.S. industries in Figure 3.6, I do not find any negative relationship between the rise in the share of services intermediate input usage and the change in the labour productivity correlations. Although all industries, except agriculture, witnessed a rise in the share of services intermediate inputs in the post-1984 period, this rise was not correlated with the industry-specific fall in the productivity correlations. All these pieces of evidence essentially refute both the composition effect and the substitution effect channels of the rise of the service sector as an explanation for the fall in labour productivity correlations.

3.5.3 Growing Share of Intangible Investment

One explanation that is provided in the literature for the drop in labour productivity correlations is the mismeasurement of output (see McGrattan and Prescott (2012), henceforth MP). The argument is that if a part of the output is not measured and if this omitted portion is more positively correlated with labour input than the measured part, then the measured labour productivity correlation can be lower than the true one. MP argues that intangible capital is one such source of mismeasurement in output, and so the increased use of intangible capital in recent years can generate countercyclical

[48] This idea of evolving input-output structure of the economy leading to switch in cyclicality of productivity can be found in Huang, Liu, and Phaneuf (2004), who explain the switch in the cyclicality of real wages in the post-War period.

labour productivity.

For the argument to hold empirically, one needs intangible investment to rise markedly around the mid-1980s. However, MP analyzes the U.S. business cycle only between 2004 and 2011. Nevertheless, it is important to corroborate whether their explanation is supported by the data when the correct time-period is considered. Specifically, I want to check if the rise in intangible investment across U.S. industries around the mid-1980s is positively correlated with the magnitude of the fall in labour productivity correlations.

McGrattan and Prescott (2012) defines intangible capital as the *"...accumulated know-how from investing in research and development, brands, and organizations, which is for the most part expensed by companies rather than capitalized."* Keeping this definition in mind, any empirical measure of intangible investment is difficult to find, but the closest one can get in available data is to look at investment in intellectual property products (IPP).[49] IPP contains research and development, computer software and databases, and other products like artistic originals.[50] While IPP investment picked up in the late 1970s and early 1980s across almost all industries, I do not find a significant correlation between the rise in IPP capital share and the drop in labour productivity correlations in the cross-section (see Figure 3.7).

3.5.4 Aggregate versus Sectoral Shocks

Aggregate productivity can be boosted through easier reallocation of factors of production across firms and industries. There is a large literature discussing the inefficiency and lost productivity due to factor misallocation (see, for example, Hsieh and Klenow (2009)). The basic intuition is that when productive factors shift to firms or industries with higher marginal products of inputs, the overall economy generates more output with the same amount of inputs, even without any technological progress. Thus, if frictions impeding the efficient allocation of resources became less important during economic downturns since the mid-1980s then measured productivity can become less procyclical. More flexible labour markets with lower frictions in hiring and firing (discussed above in Section 2.3), and deeper financial markets aiding capital movement can contribute to such countercyclical reallocation of resources in the last three decades.

In this context, Foster, Grim, and Haltiwanger (2016) find that downturns are indeed periods of accelerated reallocation that is productivity-enhancing. However, they find that the intensity of reallocation fell rather than rose for the Great Recession of 2007-08 and that the reallocation that did occur was less productivity-enhancing than in prior recessions. This casts doubt on the story of productivity-enhancing sectoral reallocation during recessions for explaining the productivity puzzle.

[49] Advertisement spending is another type of intangible investment carried out by firms. Figure 3.8 shows that there has been no sudden increase in the share of advertising expenditure in total GDP of the U.S. economy around the mid-1980s.

[50] For a detailed discussion on the measure of IPP capital used, please refer to Appendix B.4. Wang (2014) considers information and communication technology (ICT) capital as the measure of intangible capital and finds no significant correlation across U.S. industries between the drop in the correlation of TFP with primary inputs and the rise in ICT capital share.

Garin, Pries, and Sims (2018), however, differ from Foster, Grim, and Haltiwanger (2016). Using the finding in Foerster, Sarte, and Watson (2011) that sectoral reallocation shocks became more important for business cycles in the U.S. economy over the recent years, they claim that more efficient reallocation in the post-1984 period has led to less procyclical productivity. Their claim hinges on separately identifying aggregate economy-wide shocks from sector-specific shocks[51] and empirically finding that the volatility of aggregate shocks has shrunk drastically in the post-1984 era. They use monthly sectoral U.S. data from the Index of Industrial Production (IIP), which covers mostly the U.S. manufacturing sector only. I replicate the analysis using industry-level data from various sources (like the BEA, the KLEMS, and the Current Employment Statistics (CES) data) that covers the entire U.S. economy. While there is considerable heterogeneity across datasets in how much of the total variation in output and hours growth is explained by sectoral shocks, I find that, regardless of the dataset, the relative importance of sectoral shocks have increased dramatically in the post-1984 period. However, this robustness is not maintained when the number of industry classifications is large. For example, when the 31-industry classification from KLEMS dataset or the 20-industry classification from IIP data is considered, there is no clear pattern of sectoral shocks becoming more important in the later decades.[52]

Even if it is granted that reallocation of factors of production across different sectors of the economy has improved since the mid-1980s, it should be noted that sectoral measures of TFP and labour productivity already take into account the intra-industry reallocation of resources. Since a majority of U.S. industries have individually experienced a decline in procyclicality of measured productivity (as shown in Chapter 2), intra-industry reallocation across firms has likely been more important than inter-industry reallocation.[53] Therefore, while there is certainly a predominance of industry-specific shocks in the last three decades, there is reason to doubt that inter-industry reallocations have an important role to play in explaining the drop in aggregate productivity correlations.

In this context, it is also important to note that sectoral labour reallocation is also often cited as one of the major reasons for jobless recoveries after the last two recessions. This argument typically draws from the evidence in Groshen and Potter (2003) who argue that increased permanent relocation of workers from some industries to others have stalled growth in jobs. Since jobless recoveries tend to exacerbate the negative correlation between productivity and hours (although they increase the correlation between productivity and output), it is important to consider this channel of sectoral reallocation. However, Aaronson, Rissman, and Sullivan (2004) show that Groshen and Potter's findings are very sensitive to the exact period over which the measure of reallocation is computed, the dating of business cycle turning points, and the weighting of the industries. Using an alternative measure of sectoral reallocation developed by Rissman (1997) they

[51] Aggregate shock is identified as the first principal component in the data on the sectoral growth rate of output or employment. Refer to Appendix B.5 for details on the identification strategy.

[52] See Appendix B.5 for details of analysis showing the relative importance of sector-specific shocks in the U.S. economy.

[53] Wang (2014) also finds that individual industries have experienced a drop in procyclicality of procyclicality, whereas Molnarova (2020) opposes this claim.

show that reallocation of employment across industries has declined, not increased, over the past two business cycles.

To summarize, the evidence for increased inter-sectoral labour reallocation as an important explanation for either the vanishing procyclicality of labour productivity or jobless recoveries appears to be less than convincing.

3.6 Conclusion

This chapter shows that a standard New Keynesian model with endogenous effort choice in the face of costly hiring of workers can not only generate the empirically observed changes in the business cycle moments of output, employment and productivity with only a drop in hiring cost but also qualitatively match the changes in the impulse responses of these variables to technology and demand shocks. It also points out the absence of empirical evidence for plausible explanations of the productivity puzzle. In particular, it shows that neither the rise of the service sector in terms of value-added nor its use as intermediate input can explain the phenomenon. Moreover, the increased use of intangible capital, and the reduced variation of aggregate economy-wide shocks relative to sector-specific shocks that facilitate factor reallocation to more productive sectors during recessions, do not seem to have empirical validity as possible explanations for the productivity puzzle.

Important policy implications like using a more accommodative monetary policy to generate more jobs during recovery booms follow immediately from the quantitative analyses presented in this chapter. Such a policy will also help in making productivity more procyclical again. However, there are potential downsides of too much accommodation of positive technology shocks by the monetary authority as it may lose its power to fight recessions by lowering interest rates further during economic downturns. Therefore, there is a debate regarding whether maintaining countercyclical productivity, in the long run, is welfare-improving or not. Other immediate policy prescriptions like short term work (STW) policies that encourage labour hoarding by firms during recessions can also be envisaged. Giupponi and Landais (2018) show that such STW polices in Italy stabilized employment and brought small positive welfare gains during the Great Recession. Graves (2019) shows that in the United States firing taxes are more effective than hiring subsidies in stabilizing employment along the business cycle. Further research is required to shed light on these welfare implications of the productivity puzzle.

3.7 Tables

Table 3.1: Differences in Calibration between Pre- and Post-1984

Parameter	Meaning	Pre-1984	Post-1984
De-unionization			
Θ	Share of hiring cost in GDP	3%	1%
ξ	Wage Bargaining Power of Firms	0.50	0.84
Monetary Policy			
ϕ_y	Response to output gap	0.17	0.08
Shocks			
σ_a	Technology shock volatility	1.00	0.70
σ_ν	Monetary shock volatility	0.53	0.27

Table 3.2: Calibration of Time-Invariant Parameters

Parameter	Value	Calibration
β	0.99	Real risk-free annual interest rate $\simeq$ 3%
ε	10.0	Mark-up over marginal cost $\simeq$ 11%
α	0.33	Share of non-labour input in total compensation
θ_p	0.75	Calvo nominal rigidity; Galí (2011)
θ_w	0.75	Nominal wage rigidity; Galí (2011)
δ	0.10	Quarterly gross job separation rate; Shimer (2012)
ϕ_π	1.70	Taylor rule response to inflation; Smets and Wouters (2007)
$\phi_{\Delta y}$	0.20	Taylor rule response to output gap growth; Smets and Wouters (2007)
ρ	0.80	Persistence in monetary policy; Smets and Wouters (2007)
ρ_a	0.90	Persistence of technology shock; Galí (2011)
ρ_ν	0.50	Persistence of monetary policy shock; Galí (2011), Barnichon (2010)
γ	1.00	Quadratic hiring cost
ϕ	1.00	Increasing marginal disutility from effort
ψ	0.50	Additional curvature of effort in production function

Table 3.3: Changes in Business Cycle Moments due to De-unionization

Business Cycle Moments	Changes in Moments between Pre- & Post-1984 Periods			
		Model		
	Data (1)	Hiring Cost: Θ (2)	Bargaining Power: ξ (3)	De-unionization: Θ, ξ (4)
Labour Productivity Correlations				
Output: $Corr\,(y_t, lp_t)$	-0.40	-0.54	-0.08	-0.58
Employment: $Corr\,(n_t, lp_t)$	-0.51	-0.35	-0.07	-0.38
Hiring Flows: $Corr\,(h_t, lp_t)$	-0.53	-0.48	-0.04	-0.50
Relative Volatility of Employment				
Output: $s.d.\,(n_t)\,/s.d.\,(y_t)$	+46%	+35%	+6%	+38%
Conditional $\boldsymbol{Corr\,(n_t, lp_t)}$				
Technology Shock	-0.06	-0.06	-0.03	-0.07
Demand Shock	-1.24	-0.58	-0.10	-0.59
Conditional $\boldsymbol{Corr\,(y_t, lp_t)}$				
Technology Shock	+0.19	-0.08	-0.01	-0.09
Demand Shock	-1.21	-0.89	-0.11	-0.91
Real Wage Correlations				
Output: $Corr\,(y_t, w_t)$	-0.34	-0.06	-0.72	-0.41
Employment: $Corr\,(n_t, w_t)$	-0.34	-0.11	-0.77	-0.50
Labour Productivity: $Corr\,(lp_t, w_t)$	-0.10	-0.12	+0.13	+0.12

Note: To maintain comparability with the existing literature, all moments have been calculated on quarterly Hodrick and Prescott (1997) filtered variables for both data and model-simulated series. Column (1) reports the empirically observed changes in the business cycle moments between the pre- and post-1984 periods. Column (2) reports the changes in the model-implied moments when only the hiring cost parameter Θ is allowed to drop from 3% to 1%. Similarly, column (3) allows only the wage bargaining power parameter ξ to increase from 0.50 to 0.84. Column (4) combines the two parameter changes in columns (2) and (3).

Table 3.4: Changes in Business Cycle Moments between Pre- and Post-1984

Business Cycle Moments	**Changes in Moments between Pre- & Post-1984 Periods**			
		Model		
	Data (1)	De-unionization (2)	Monetary Policy: ϕ_y (3)	Shocks: σ_a,σ_ν (4)
Labour Productivity Correlations				
Output: $Corr\,(y_t, lp_t)$	-0.40	-0.58	-0.01	+0.09
Employment: $Corr\,(n_t, lp_t)$	-0.51	-0.38	+0.03	-0.04
Hiring Flows: $Corr\,(h_t, lp_t)$	-0.53	-0.50	+0.05	-0.10
Relative Volatility of Employment				
Output: $s.d.\,(n_t)\,/s.d.\,(y_t)$	+46%	+38%	-1%	-2%
Conditional $Corr\,(n_t, lp_t)$				
Technology Shock	-0.06	-0.07	+0.05	0.00
Demand Shock	-1.24	-0.59	-0.01	0.00
Conditional $Corr\,(y_t, lp_t)$				
Technology Shock	+0.19	-0.09	+0.04	0.00
Demand Shock	-1.21	-0.91	-0.01	0.00
Real Wage Correlations				
Output: $Corr\,(y_t, w_t)$	-0.34	-0.41	+0.02	-0.00
Employment: $Corr\,(n_t, w_t)$	-0.34	-0.50	+0.09	-0.16
Labour Productivity: $Corr\,(lp_t, w_t)$	-0.10	+0.12	-0.08	+0.15

Note: To maintain comparability with the existing literature, all moments have been calculated on quarterly Hodrick-Prescott-filtered variables for both data and model-simulated series. Column (2) refers to the total effect of de-unionization by allowing the parameters Θ and ξ to change, same as column (4) in Table 3.3. Column (3) allows only the Taylor rule parameter ϕ_y to drop from 0.17 to 0.08. Similarly, column (4) corresponds to the changes in model-implied moments when only the volatilities of the shocks are allowed to decrease in the post-1984 period according to the calibration in Table 3.1.

Table 3.5: Robustness to Choice of γ

Business Cycle Moments	**Changes in Model Moments**		
	$\gamma = 0.6$ (1)	Baseline, $\gamma = 1$ (2)	$\gamma = 2.4$ (3)
Labour Productivity Correlations			
Output: $Corr\,(y_t, lp_t)$	-0.44	-0.58	-0.64
Employment: $Corr\,(n_t, lp_t)$	-0.27	-0.38	-0.49
Hiring Flows: $Corr\,(h_t, lp_t)$	-0.44	-0.50	-0.45
Relative Volatility of Employment			
Output: $s.d.\,(n_t)\,/s.d.\,(y_t)$	+28%	+38%	+66%
Conditional $\boldsymbol{Corr\,(n_t, lp_t)}$			
Technology Shock	-0.07	-0.07	-0.05
Demand Shock	-0.37	-0.59	-0.74
Conditional $\boldsymbol{Corr\,(y_t, lp_t)}$			
Technology Shock	-0.07	-0.09	-0.08
Demand Shock	-0.63	-0.91	-0.97
Real Wage Correlations			
Output: $Corr\,(y_t, w_t)$	-0.29	-0.41	-0.65
Employment: $Corr\,(n_t, w_t)$	-0.37	-0.50	-0.76
Labour Productivity: $Corr\,(lp_t, w_t)$	+0.12	+0.12	+0.07

Note: Columns (1) through (3) report changes in business cycle moments between pre- and post-1984 periods for alternative values of parameter, γ, denoting the degree of convexity of the hiring cost function. All other parameters in the model are fixed at the calibration values used in column (4) of Table 3.3, which corresponds to the total effect of de-unionization. To maintain comparability with the existing literature, all moments have been calculated on quarterly Hodrick-Prescott-filtered variables.

Table 3.6: Robustness to Choice of ϕ

Business Cycle Moments	**Changes in Model Moments**		
	$\phi = 0.5$	Baseline, $\phi = 1$	$\phi = 1.5$
	(1)	(2)	(3)
Labour Productivity Correlations			
Output: $Corr\,(y_t, lp_t)$	-0.62	-0.58	-0.54
Employment: $Corr\,(n_t, lp_t)$	-0.42	-0.38	-0.35
Hiring Flows: $Corr\,(h_t, lp_t)$	-0.52	-0.50	-0.48
Relative Volatility of Employment			
Output: $s.d.\,(n_t)\,/s.d.\,(y_t)$	+43%	+38%	+34%
Conditional $\boldsymbol{Corr\,(n_t, lp_t)}$			
Technology Shock	-0.06	-0.07	-0.07
Demand Shock	-0.65	-0.59	-0.54
Conditional $\boldsymbol{Corr\,(y_t, lp_t)}$			
Technology Shock	-0.09	-0.09	-0.09
Demand Shock	-0.97	-0.91	-0.85
Real Wage Correlations			
Output: $Corr\,(y_t, w_t)$	-0.38	-0.41	-0.42
Employment: $Corr\,(n_t, w_t)$	-0.48	-0.50	-0.50
Labour Productivity: $Corr\,(lp_t, w_t)$	+0.08	+0.12	+0.15

Note: Columns (1) through (3) report changes in business cycle moments between pre- and post-1984 periods for alternative values of parameter, ϕ, denoting the degree of increasing marginal disutility from exerting more effort. All other parameters in the model are fixed at the calibration values used in column (4) of Table 3.3, which corresponds to the total effect of de-unionization. To maintain comparability with the existing literature, all moments have been calculated on quarterly Hodrick-Prescott-filtered variables.

Table 3.7: Robustness to Choice of ψ

Business Cycle Moments	**Changes in Model Moments**		
	$\psi = 0.10$	$\psi = 0.25$	Baseline, $\psi = 0.50$
	(1)	(2)	(3)
Labour Productivity Correlations			
Output: $Corr\,(y_t, lp_t)$	-0.30	-0.46	-0.58
Employment: $Corr\,(n_t, lp_t)$	-0.19	-0.29	-0.38
Hiring Flows: $Corr\,(h_t, lp_t)$	-0.29	-0.43	-0.50
Relative Volatility of Employment			
Output: $s.d.\,(n_t)\,/s.d.\,(y_t)$	+19%	+28%	+38%
Conditional $\boldsymbol{Corr\,(n_t, lp_t)}$			
Technology Shock	-0.11	-0.08	-0.07
Demand Shock	-0.20	-0.43	-0.59
Conditional $\boldsymbol{Corr\,(y_t, lp_t)}$			
Technology Shock	-0.02	-0.08	-0.09
Demand Shock	-0.33	-0.69	-0.91
Real Wage Correlations			
Output: $Corr\,(y_t, w_t)$	-0.44	-0.43	-0.41
Employment: $Corr\,(n_t, w_t)$	-0.53	-0.51	-0.50
Labour Productivity: $Corr\,(lp_t, w_t)$	+0.27	+0.19	+0.12

Note: Columns (1) through (3) report changes in business cycle moments between pre- and post-1984 periods for alternative values of parameter, ψ, denoting the additional curvature for effort in the production function. Given $\alpha = 0.33$ in the baseline calibration, $\psi \in (0, 0.50]$ to ensure non-increasing returns to scale under perfect competition among intermediate goods firms. All other parameters in the model are fixed at the calibration values used in column (4) of Table 3.3, which corresponds to the total effect of de-unionization. To maintain comparability with the existing literature, all moments have been calculated on quarterly Hodrick-Prescott-filtered variables.

Table 3.8: Robustness to Changes in Nominal Rigidities

Business Cycle Moments		**Changes in Moments due to De-unionization** Model	
	Data (1)	**No change in rigidity** (2)	**Changes in θ_p & θ_w** (3)
Labour Productivity Correlations			
Output: $Corr\left(y_t, lp_t\right)$	-0.40	-0.58	-0.87
Employment: $Corr\left(n_t, lp_t\right)$	-0.51	-0.38	-0.43
Hiring Flows: $Corr\left(h_t, lp_t\right)$	-0.53	-0.50	-0.33
Relative Volatility of Employment			
Output: $s.d.\left(n_t\right)/s.d.\left(y_t\right)$	+46%	+38%	+73%
Conditional $\boldsymbol{Corr\left(n_t, lp_t\right)}$			
Technology Shock	-0.06	-0.07	-0.11
Demand Shock	-1.24	-0.59	-0.55
Conditional $\boldsymbol{Corr\left(y_t, lp_t\right)}$			
Technology Shock	+0.19	-0.09	-0.17
Demand Shock	-1.21	-0.91	-0.90
Real Wage Correlations			
Output: $Corr\left(y_t, w_t\right)$	-0.34	-0.41	-0.20
Employment: $Corr\left(n_t, w_t\right)$	-0.34	-0.50	-0.20
Labour Productivity: $Corr\left(lp_t, w_t\right)$	-0.10	+0.12	-0.24

Note: Column (2) corresponds to $\theta_p = \theta_w = 0.75$ for both periods as in the baseline calibration. Column (3) corresponds to changing θ_p from 0.55 to 0.73, and θ_w from 0.65 to 0.74 between the pre- and post-1984 periods, along with the changes in Θ and ξ like in column (4) of Table 3.3.

Table 3.9: Labour Productivity Correlations in Manufacturing & Services

	With Output			**With Hours**		
Sector	Pre-1983	Post-1984	Change	Pre-1983	Post-1984	Change
Manufacturing	0.63	0.40	-0.23	-0.04	-0.30	-0.26
Services	0.68	0.48	-0.20	-0.10	-0.59	-0.49

Note: Data is sourced from annual KLEMS dataset between 1947 and 2010 by aggregating industry-level non-additive chained indices according to the cyclical expansion method developed in Cassing (1996). Results are robust to using annual sectoral dataset from BEA, compiled by Herrendorf, Herrington, and Valentinyi (2015).

3.8 Figures

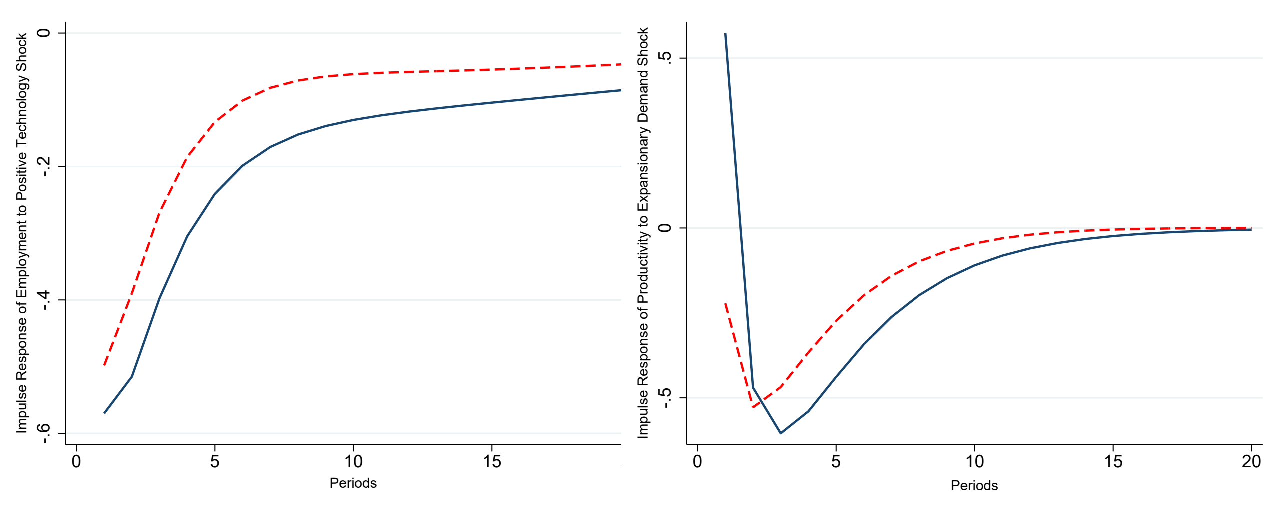

(a) Technology Shock: Employment (b) Demand Shock: Labour Productivity

Figure 3.1: Model-implied Impulse Responses to Technology and Demand Shocks

Note: Model-generated Impluse Response Functions (IRF) for the pre-1984 period are in blue, and the post-1984 IRF's are in red dashed lines. Pre- and post-1984 calibrations of parameters correspond to all parameter changes listed in Table 3.1.

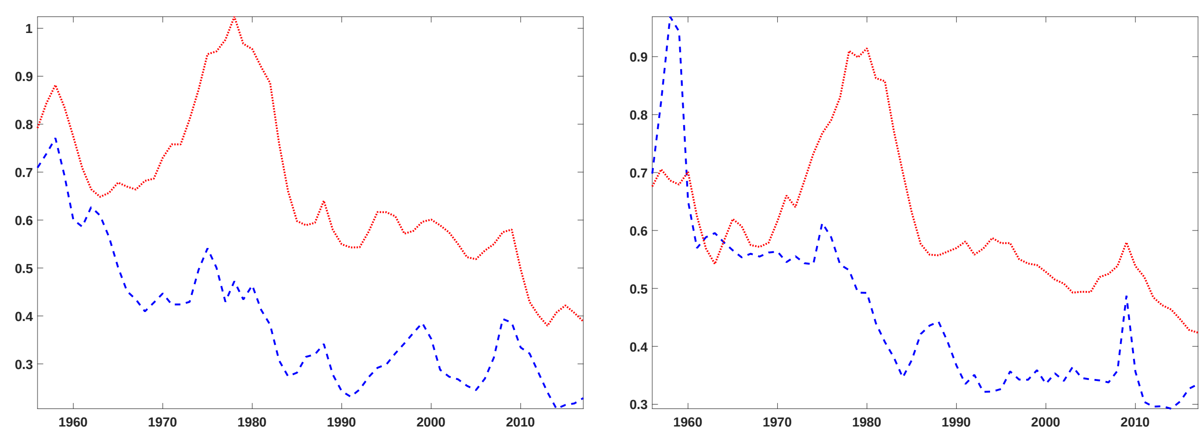

(a) Labour Productivity (b) Total Factor Productivity

Figure 3.2: Conditional Volatility of Hours

Note: Time-varying standard deviations of per capita hours, conditional on technology shock (blue dashed line) and demand shock (red dotted line). Measures of productivity are different in the two panels. Panels (a) and (b) use growth rates of labour productivity and TFP respectively, along with per capita hours as the two variables in the SVAR.

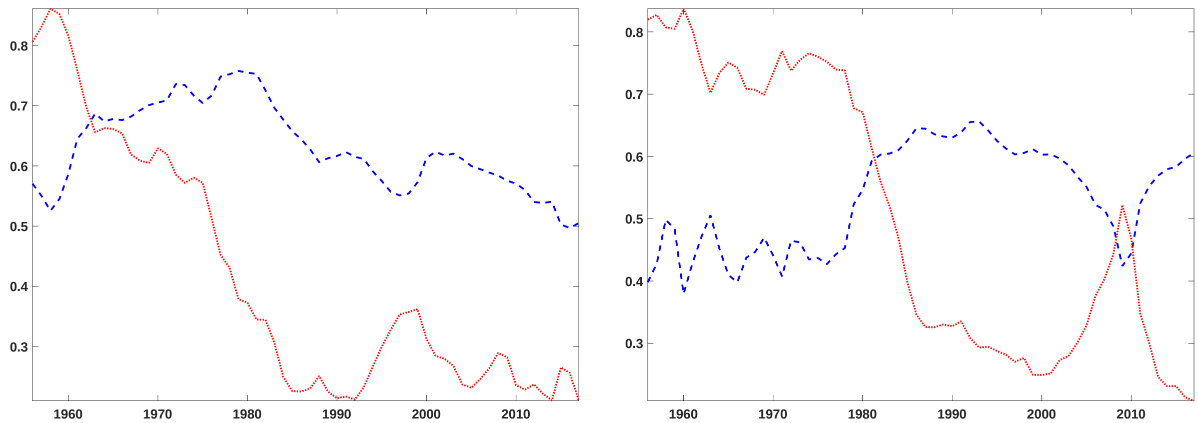

(a) Labour Productivity

(b) Total Factor Productivity

Figure 3.3: Conditional Volatility of Productivity

Note: Time-varying standard deviations of productivity, conditional on technology shock (blue dashed line) and demand shock (red dotted line). Measures of productivity are different in the two panels. Panels (a) and (b) use growth rates of labour productivity and TFP respectively, along with per capita hours as the two variables in the SVAR.

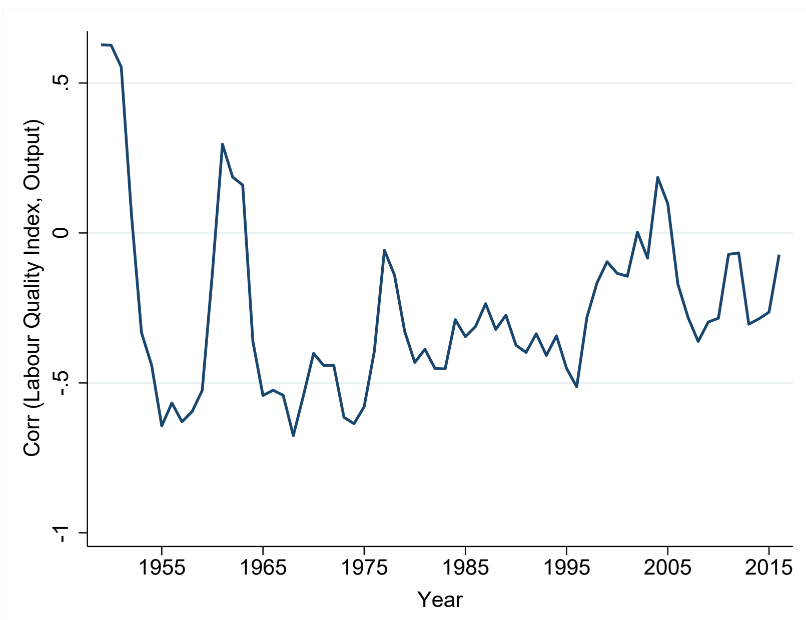

Figure 3.4: Correlation of Labour Quality Index with Output

Note: Labour Quality Index is sourced from Fernald (2014). From 1979 onwards it is the one constructed by Aaronson and Sullivan (2001). Pre-1979 data is interpolated annual BLS Multi Factor Productivity estimate of labour composition using method in Denton (1971).

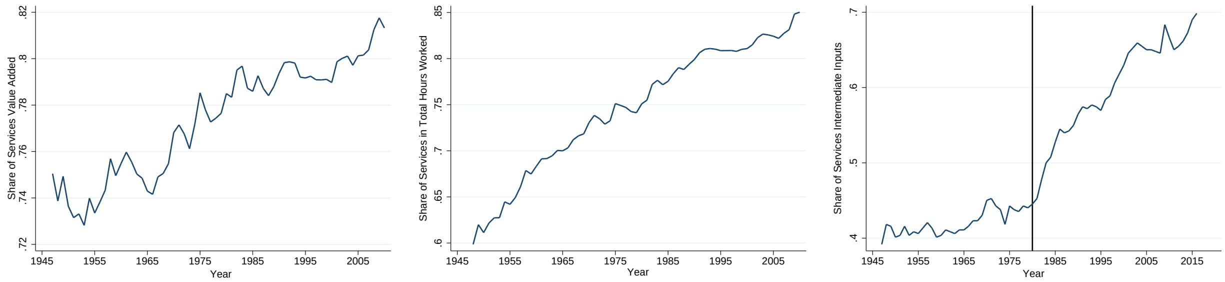

(a) Output Share (b) Hours Share (c) Intermediate Inputs Share

Figure 3.5: Share of Services in the U.S. (1947-2016)

Note: Data for Panels (a) and (b) is sourced from KLEMS annual dataset. Data for Panel (c) is sourced from the annual input-output matrices published by BEA.

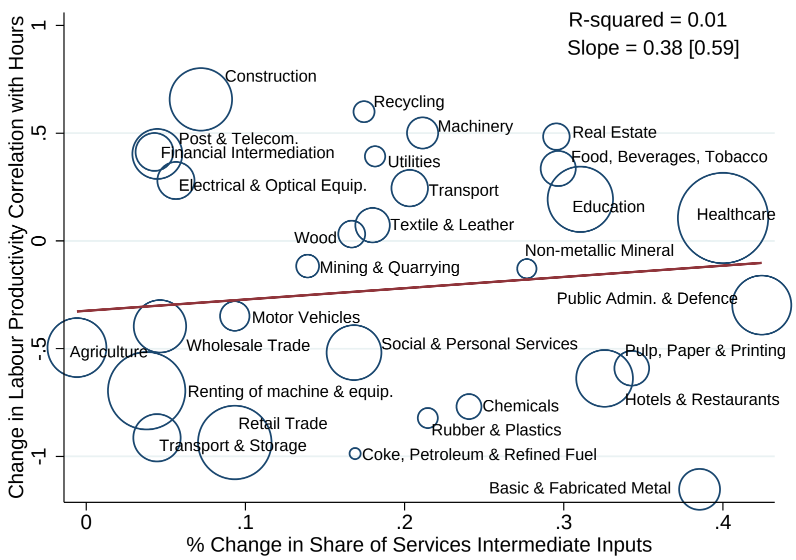

Figure 3.6: Changes in Share of Services Intermediate Input & Labour Productivity Correlation

Note: Data for labour productivity correlations and share of services intermediate inputs at the industry-level is sourced from the annual KLEMS dataset. Time-changes refer to the difference between the average values in the post-1984 (1984-2010) and the pre-1984 period (1969-1983). Regression is weighted by the time-average of total hours worked in each industry, depicted by the size of the bubbles. The p-value of the estimated slope is reported in parentheses. The Baxter and King (1999) bandpass filter between 2 and 8 years has been used to extract the cyclical component of the variables. Result is robust to using other filters and time-horizons.

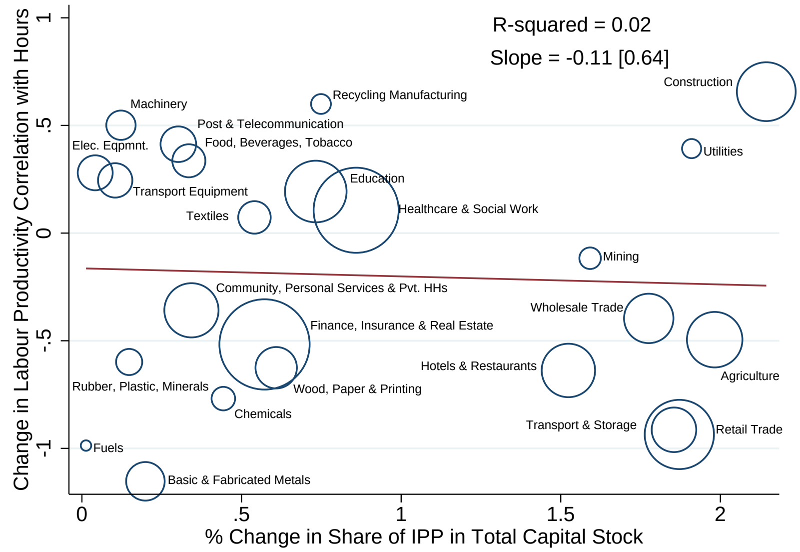

Figure 3.7: Changes in Share of IPP in Total Capital Stock & Labour Productivity Correlation

Note: Data for labour productivity correlations at the industry-level is sourced from the annual KLEMS dataset, and that for the IPP capital share is sourced from BEA. Industry codes from the two datasets were matched to create a consistent set of 24 U.S. industries. Time-changes refer to the difference between the average values in the post-1984 (1984-2010) and the pre-1984 period (1969-1983). Regression is weighted by the time-average of industry employment, depicted by the size of the bubbles. The p-value of the estimated slope is reported in parentheses. The Baxter and King (1999) bandpass filter between 2 and 8 years has been used to extract the cyclical component of the variables. Result is robust to using other filters and time-horizons.

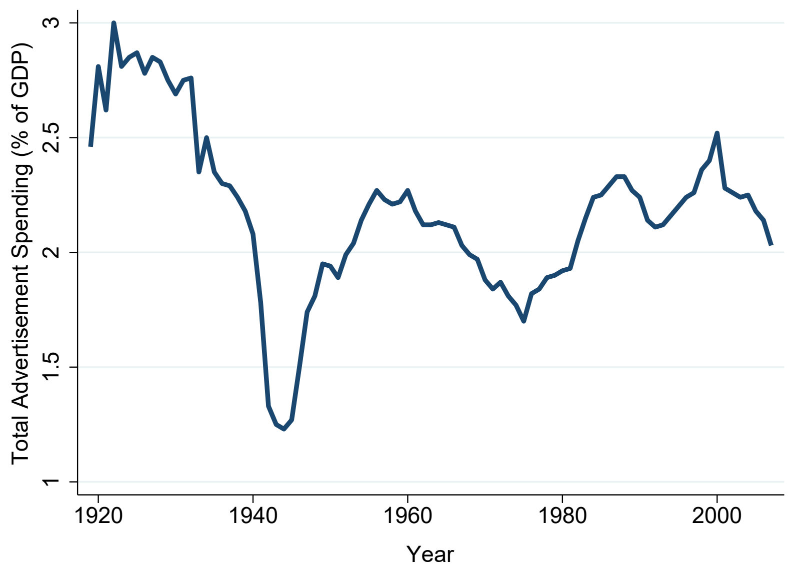

Figure 3.8: Total Advertisement Spending as a Share of GDP in the U.S. (1919-2007)

Note: Data on advertisement spending is maintained by Douglas A. Galbi at www.galbithink.org/ad-spending.htm

Chapter 4

Consumption and Income Inequality across Generations

4.1 Introduction

How do parents impact the economic outcomes of their offspring? The current chapter tries to provide a comprehensive understanding of the diverse channels of parental influence, namely, through labour earnings, other sources of income such as inter-vivos and bequest transfers and spousal income, and through consumption. It also quantifies how much of the observed economic inequality in a particular generation is attributable to intra-family inter-generational linkages.

We argue that since the various channels of transmission of parental characteristics are interlinked (e.g., spending on a child's education may be a substitute for income transfers in adulthood, etc.), it is important to consider the *joint* evolution of these different income sources and consumption expenditure. Crucially, this joint estimation strategy involving a system of moment conditions not only enables simultaneous estimation of cross-sectional inequality parameters and inter-generational persistence parameters but also allows for potentially non-trivial covariances among income sources and consumption. We find that intergenerational persistence is highest for lifetime earnings, with an elasticity around 0.23. In contrast, the inter-generational elasticity for other income is only 0.10 and mostly reflects persistence in spousal earnings. That is, men tend to marry women who have similar economic outcomes as their mothers (as in Fernandez, Fogli, and Olivetti (2004)). Interestingly, other income has little effect on children's earnings with an elasticity of 0.06 but we find evidence that higher parental earnings are associated with higher unearned income among kids, with a cross-elasticity of 0.21 in the baseline specification. Restricting these latter cross-elasticities to zero leads to under-estimates of the importance of parents for consumption inequality. We estimate a significant direct consumption pass-through, albeit a little weaker than the pass-through of earnings. Of course, consumption persistence operates also indirectly through income channels. Finally, we also find that persistence in observable characteristics accounts for most of the intergenerational pass-through. Taken together, our estimates of the intergenerational pass-through are consistent with the view that persistence is largely driven by associations in lifetime earnings of both spouses as well as tastes and preferences in consumption, with educational attainment playing a crucial role.[54]

[54] See Landersø and Heckman (2017) for evidence on the importance of education for intergenerational persistence. For a discussion of causal effects of parental education and income see, among others, Carneiro and Heckman (2003); Oreopoulos, Page, and Stevens (2006); Belley and Lochner (2007); Black and Devereux (2011); Holmlund, Lindahl,

The central question that we address in the debate on the role of family background for life-cycle outcomes is whether observed within-generation inequality would be much different if heterogeneity among parents were removed. Our model delivers a transparent setting to perform inequality accounting exercises and quantify the contribution of parental factors. These exercises consistently indicate that idiosyncratic heterogeneity over the life cycle, rather than family background, accounts for the bulk of cross-sectional dispersion in earnings, income and expenditure. The largest impact of parental factors is on consumption inequality, as our baseline estimates imply that roughly one-third of within-generation consumption inequality can be attributed to family characteristics. Further examination shows that the relatively larger role of family heterogeneity on consumption follows from the interaction of (i) cross-sectional insurance reducing the impact of idiosyncratic income risk on expenditures, and (ii) the direct and indirect parental influences that are reflected in consumption choices, notably the intergenerational transmission of saving propensities (higher savings rates for richer families) and marital sorting.

We show that the historical evolution of consumption and earnings inequality is consistent with stable intergenerational pass-through coefficients. This result helps reconcile the somewhat puzzling observation of fairly stable intergenerational persistence (Hertz, 2007; Lee and Solon, 2009) in the face of growing inequality (Heathcote, Perri, and Violante, 2010; Attanasio and Pistaferri, 2016). Our estimates show that growing parental disparities would not, all else equal, be sufficient to trigger significantly higher future inequality in the absence of much stronger inter-generational elasticities.[55]

The rest of the chapter is organized as follows. Section 4.2 outlines our benchmark consumption model with intergenerational linkages. Section 4.3 discusses the identification of model parameters, outlines the estimation approach and describes the data from the Panel Study of Income Dynamics (PSID) that is used for estimation. Baseline estimates are discussed in Section 4.4, where we also report the results of the counterfactual analysis of cross-sectional inequality. In Section 4.5 we explore the implications of our estimates for the evolution of cross-sectional inequality. Section 4.6 presents various robustness checks, including alternative approaches to the measurement of consumption expenditures and different model specifications. Section 4.7 concludes.

4.2 A Model of Intergenerational Inequality

We develop an estimable consumption model of heterogeneous and intergenerationally linked households. The model features multiple parent-child linkages and is designed to examine the joint behaviour of earnings, other income and expenditures.

To motivate these linkages, we begin by establishing stylized facts about the evolution of intra-family persistence in the U.S. over recent decades. In Appendix C.3.1 we report reduced-form

and Plug (2011); Lefgren, Sims, and Lindquist (2012); Lee, Roys, and Seshadri (2014). Carneiro, Lopez Garcia, Salvanes, and Tominey (2015) show that the timing of parental income affects children education outcomes with the advantage of having income occur in late adolescence rather than early childhood.

[55] See Cordoba, Liu, and Ripoll (2016) for a dynastic model of long-run inequality and social mobility with endogenous fertility.

estimates of the intergenerational pass-through of earnings and consumption since 1990, obtained using the method popularized by Lee and Solon (2009) in their analysis of the gender-specific evolution of earnings persistence. Like those authors, we find little or no evidence of changes in the intergenerational elasticity of labour earnings, with similar patterns holding for expenditures.[56] To corroborate this evidence, we also compute mobility matrices and intergenerational flows across quartiles of the distributions of earnings and expenditures.[57] This analysis, shown in Appendix C.3.2, emphasizes that persistence is more intense at the tails of the distribution and that the inter-generational pass-through was remarkably stable over the past decades. These findings are consistent with the evidence in Chetty, Hendren, Kline, Saez, and Turner (2014), who examine large administrative U.S. earning records and conclude that measures of "...intergenerational mobility have remained extremely stable for the 1971-1993 birth cohorts". For these reasons, we maintain the assumption of stationarity in the baseline analysis. However, among the robustness checks of Section 4.6, we explore potential cross-cohort differences in the cross-generation pass-through parameters and the variances of the idiosyncratic risk processes.

The building blocks of our analysis are the time series processes for earned and other income of parents and children, and a mechanism mapping them into distributions of family outcomes. Each household optimally chooses per-period consumption expenditures to maximize discounted expected utility subject to a budget constraint.[58] Households receive income from labour earnings of the head (the husband for couples in the PSID), as well as from 'other income' that includes transfer income and earnings of the spouse. Households also accrue income from asset returns, which implicitly depend on their consumption and saving decisions. We allow for each of labour earnings and other income to be individually linked across generations. Moreover, we let the income from asset returns be part of a residual consumption shifter that can also be linked across generations.[59] While consumption is modelled as an optimal state-dependent choice, due to insufficient data we do not impose the adding up required to satisfy a lifetime budget constraint. This would require more frequent and accurate data on the full set of expenditures, flows of asset income, inter-vivos transfers and bequests.

Earnings and Other Income. We denote a time-period (a year) by t. Parent and child are identified by superscripts p and k. A parent-child pair is denoted by the family subscript f. Head's earnings, other income and consumption expenditures (all logged) are denoted by e, n and c, respectively.

Our baseline specification of the parents' earning process has a canonical permanent-transitory form. Specifically, it features a permanent individual fixed effect and an additive transitory shock

[56] This is despite growing income and consumption inequality over this period (e.g. Aguiar and Bils, 2015).

[57] Mobility matrices deliver the conditional probability of a child being placed in a certain quartile of the distribution given the quartile of his/her family.

[58] This differs from approaches in which consumption expenditures were modelled as a stand-alone exogenous process.

[59] The PSID does not report consistent wealth information before 1998. Therefore, we subsume unobserved income from wealth within the residual consumption shifter terms.

component. The fixed effect measures lifetime average earnings and should not be thought of as pre-determined at the start of an individual's life. Rather, as shocks accrue over the working life, earnings in any period can exceed, or fall below, the lifetime average. Hence, we use the fixed effect representation as a way of summarising the lifetime resources available to a particular generation. In Section 4.6.3 we also consider robustness to an alternative model specification that focuses on growth rates. The latter allows for period-specific permanent innovations that are correlated across generations. We find no evidence in support of this alternative specification of parent-child linkages.

In any year t the parent in family f has earnings $e^p_{f,t}$ equal to the sum of an individual fixed effect component, $\bar{e}^p_f$, and an independent mean zero transitory shock, $\zeta^p_{f,t}$, with variance $\sigma^2_{\zeta^p}$. Similarly, the process for other income, $n^p_{f,t}$, comprises a permanent component, $\bar{n}^p_f$, and a transitory mean zero component, $u^p_{f,t}$, with variance $\sigma^2_{u^p}$.

$$e^p_{f,t} = \bar{e}^p_f + \zeta^p_{f,t} \tag{4.1}$$
$$n^p_{f,t} = \bar{n}^p_f + u^p_{f,t}. \tag{4.2}$$

The income processes of children are also assumed to have a canonical permanent-transitory structure; that is, $e^k_{f,t} = \bar{e}^k_f + \zeta^k_{f,t}$ and $n^k_{f,t} = \bar{n}^k_f + u^k_{f,t}$ (where $\zeta^k_{f,t}$ and $u^k_{f,t}$ are mean zero i.i.d. innovations with variances $\sigma^2_{\zeta^k}$ and $\sigma^2_{u^k}$ respectively). Fixed effects are partly determined by parental permanent components but also depend on idiosyncratic random variables that are independent of parents. For the children of family f this structure results in the following income components:

$$e^k_{f,t} = \underbrace{\gamma\bar{e}^p_f + \theta\bar{n}^p_f + \delta^k_f}_{\bar{e}^k_f} + \zeta^k_{f,t} \tag{4.3}$$
$$n^k_{f,t} = \underbrace{\rho\bar{n}^p_f + \lambda\bar{e}^p_f + \varepsilon^k_f}_{\bar{n}^k_f} + u^k_{f,t} \tag{4.4}$$

where ε^k_f and δ^k_f are idiosyncratic permanent shocks with variances $\sigma^2_{\varepsilon^k}$ and $\sigma^2_{\delta^k}$, respectively.

We allow for the most general dependence structure across generations: not only is there a direct channel from parental earnings to child earnings, and a direct channel from other income of parents to other income of children, but there are also cross effects. Parental earnings can affect other income of children, and parental other income can affect earnings of children; that is, higher parental lifetime earnings can influence child earnings through the persistence parameter γ, but can also change a child's other income as captured by the parameter λ.

Consumption. With the income processes in place, we solve the dynamic life-cycle problem that delivers consumption choices. When a household makes consumption decisions, it has knowledge of its own permanent income, but does not know the value of future income shocks. Our approach does not specify an altruistic motive and is agnostic about it.[60] The consumption problem of a member of family f is given by:

$$\begin{aligned}\max_{\{C_{f,k}\}_{k=t}^{T}} \quad & \mathbb{E}_t \quad \sum_{j=0}^{T-t} \beta^j u(C_{f,t+j}) \\ & s.t. \\ A_{f,t+1} \quad = \quad & (1+r)\left(A_{f,t}+E_{f,t}+N_{f,t}-C_{f,t}\right),\end{aligned} \tag{4.5}$$

where β is the discount factor, r is the real interest rate, $A_{f,t}$ is assets at the start of the period, $E_{f,t}$ is the value of the household head's labour earnings, and $N_{f,t}$ is the value of other income.[61] Consumption at time t is the annuity value of lifetime resources. The approximate log-consumption process[62] for a parent can be represented as,

$$c^p_{f,t} = q^p_{f,t} + \bar{e}^p_f + \bar{n}^p_f + \alpha\left(r\right)\left(u^p_{f,t} + \zeta^p_{f,t}\right).$$

The term $\alpha\left(r\right)$ is an annuitization factor which tends to $r/(1+r)$ as the time horizon becomes larger. The variable $q^p_{f,t}$ denotes an idiosyncratic consumption shifter, subsuming unobserved income from savings as well as possible heterogeneity in preferences over the timing of consumption. Like other shifters of consumption, $q^p_{f,t}$ comprises both a permanent and a transitory component so that $q^p_{f,t} = \bar{q}^p_f + \upsilon^p_{f,t}$. Combining these processes, the log-consumption of the parent can be written as:

$$c^p_{f,t} = \bar{q}^p_f + \bar{e}^p_f + \bar{n}^p_f + \upsilon^p_{f,t} + \alpha(r)\left(u^p_{f,t} + \zeta^p_{f,t}\right) \tag{4.6}$$

and analogously for the child. Parents affect the consumption of their children through family persistence in both earnings and other income, as described in (4.3) and (4.4). In addition, we allow for a direct transmission channel through the consumption shifter $\bar{q}^k_f$, which comprises an inherited component and a child-specific component: $\bar{q}^k_f = \phi\bar{q}^p_f + \psi^k_f$. Substituting the intra-family transmission mechanisms into the log-consumption process, we obtain:

$$\begin{aligned}c^k_{f,t} \quad = \quad & \phi\bar{q}^p_f + \left(\gamma+\lambda\right)\bar{e}^p_f + \left(\rho+\theta\right)\bar{n}^p_f \\ + \quad & \varepsilon^k_f + \psi^k_f + \delta^k_f + \upsilon^k_{f,t} + \alpha(r)\left(u^k_{f,t} + \zeta^k_{f,t}\right).\end{aligned} \tag{4.7}$$

There are, therefore, three ways in which parents can affect the consumption process of their children: (i) the earnings potential channel; (ii) the transfers and other income channel; and (iii)

[60]For an analysis of altruistically linked households using PSID data see Altonji, Hayashi, and Kotlikoff (1992, 1997).

[61]In our baseline estimation we define other income as the sum of spousal earnings and total transfer income of husband and wife. In robustness checks we split other income into its two components and assess which one accounts for most of the estimated persistence.

[62]Appendix C.1 reports analytical solutions obtained by (i) assuming a quadratic utility function or (ii) a first-order Taylor approximation of the Euler equation under CRRA utility.

inherited consumption shifters.

4.2.1 Cross-sectional Insurance and Intergenerational Smoothing

The presence of an intergenerational correlation in the consumption shifter $q_{f,t}$ reflects the accrual of different family influences. In particular, heterogeneity in $q_{f,t}$ may capture family-specific consumption preferences that shape saving behaviour. As we show in Appendix C.1, linear approximations of the Euler equation for general concave utility functions (say, CRRA) lead to omitted higher-order preference terms being loaded onto the unobserved $q_{f,t}$ shifter.[63] Accounting for the co-dependence between consumption propensities and income turns out to be quantitatively important (see Alan, Browning, and Ejrnæs, 2018).[64] In the estimation, we find evidence of strong negative covariance between consumption shifters $\bar{q}_f$ and measures of income. One interpretation of this negative correlation is that households with higher income tend to save proportionally more. This behaviour acts as a force towards reducing the cross-sectional dispersion of expenditures.

The model suggests that two competing mechanisms shape the distribution of consumption. First, a dampening effect whereby the negative correlation of consumption and saving propensities with income compresses the cross-sectional variance of household expenditures (see Blundell, Pistaferri, and Preston, 2008; Kaplan and Violante, 2010)). Second, an intra-family smoothing mechanism, whereby parents attempt to equalize marginal utilities of family members across generations. This intergenerational smoothing has the effect of inflating inequality in both consumption and income in the generation of children. Both mechanisms find support in our empirical analysis.

Breaking Down Inequality. Equations (4.1) to (4.7) specify the complete set of conditions that characterize intergenerational dependence in this economy, linking earned income, non-earned income and consumption across generations. Once estimated, these relationships can be used to characterise inequality among parents and children, as well as to highlight how parental heterogeneity translates into inequality among their children. Equations (4.1), (4.2) and (4.6) describe the processes (in levels) for parents and can be mapped into cross-sectional variances:

$$\operatorname{Var}\left(e_f^p\right) = \sigma^2_{\bar{e}^p} \tag{4.8}$$

$$\operatorname{Var}\left(n_f^p\right) = \sigma^2_{\bar{n}^p} \tag{4.9}$$

$$\operatorname{Var}\left(c_f^p\right) = \sigma^2_{\bar{q}^p} + \sigma^2_{\bar{e}^p} + \sigma^2_{\bar{n}^p} + 2\left(\sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right). \tag{4.10}$$

The latter equations highlight how consumption inequality among parents depends not only on inequality in earnings and other income, but also on their covariances. To the extent that insurance

[63] Higher order preference terms may co-move with earnings $\bar{e}_f$ and with other income $\bar{n}_f$. For example, if individuals with lower permanent income are credit-constrained, a precautionary saving motive might generate a negative correlation between $q_{f,t}$ and permanent income (see Caballe, 2016).

[64] Alan, Browning, and Ejrnæs (2018) show that consumption responses to income shocks are heterogeneous and exhibit significant co-dependence with preference shifters and idiosyncratic properties of income.

implies that other income is negatively correlated with earnings, then consumption inequality may be lower than earnings inequality.

Similarly, equations (4.3), (4.4) and (4.7) describe the key processes (in levels) for children and how inequality among children depends on inequality among parents:

$$
\begin{aligned}
\operatorname{Var}\left(e_f^k\right) &= \gamma^2\sigma^2_{\bar{e}^p} + \theta^2\sigma^2_{\bar{n}^p} + 2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p} + \sigma^2_{\delta^k} && (4.11)\\
\operatorname{Var}\left(n_f^k\right) &= \rho^2\sigma^2_{\bar{n}^p} + \lambda^2\sigma^2_{\bar{e}^p} + 2\rho\lambda\sigma_{\bar{e}^p,\bar{n}^p} + \sigma^2_{\varepsilon^k} && (4.12)\\
\operatorname{Var}\left(c_f^k\right) &= \phi^2\sigma^2_{\bar{q}^p} + (\gamma+\lambda)^2\sigma^2_{\bar{e}^p} + (\rho+\theta)^2\sigma^2_{\bar{n}^p} \\
&\quad +2\left[(\gamma+\lambda)\,\phi\sigma_{\bar{q}^p,\bar{e}^p} + (\rho+\theta)\,\phi\sigma_{\bar{q}^p,\bar{n}^p} + (\rho+\theta)(\gamma+\lambda)\,\sigma_{\bar{e}^p,\bar{n}^p}\right] \\
&\quad +\sigma^2_{\varepsilon^k} + \sigma^2_{\psi^k} + \sigma^2_{\delta^k} + 2\left[\sigma_{\psi^k,\varepsilon^k} + \sigma_{\psi^k,\delta^k} + \sigma_{\delta^k,\varepsilon^k}\right]. && (4.13)
\end{aligned}
$$

Earnings inequality among children changes with (i) the magnitude of earnings inequality among parents ($\sigma^2_{\bar{e}^p}$) and (ii) the intensity of the intergenerational pass-through (γ). It is, therefore, clear that the pass-through parameter alone is not sufficient to determine how much parents matter for inequality in subsequent generations. For consumption, the first two rows of equation (4.13) describe how parental heterogeneity drives differences among their offspring: the first row captures the direct effects of inequality among parents being transmitted into inequality among children; the second row describes the covariances which may offset the direct effects. Finally, the last row captures the drivers of inequality among children that are independent of parents.

4.3 Identification, Estimation and Data

To identify and estimate the drivers of inequality among children, as captured in equations (4.11), (4.12) and (4.13), we focus on lifetime inequality and abstract from transitory components of income and consumption. We revisit the role of transitory components in Section 4.6 where we document the robustness of baseline estimates to the inclusion of yearly variation induced by transitory shocks.[65]

4.3.1 Identification

Identification proceeds in three steps. First, we use cross-sectional moments for parents and identify variances and covariances between their sources of income and consumption. Second, we use these estimates, and inter-generational covariances, to recover parent-child persistence parameters. Lastly, information from the previous two steps is used alongside second moments from the cross-section of children to identify specific components driving inequality among children.

[65] Transitory shocks cannot be identified separately from classical measurement error in the observed variables. Bound, Brown, Duncan, and Rodgers (1994) provide estimates of the proportion of variance in observed earnings that can be attributed to measurement error. In the baseline specification we use time-averaged observations for a cross-section of individuals to mitigate concerns about classical measurement errors.

Cross-sectional Parameters for Parents. To identify parental dispersion parameters we use equations (4.8), (4.9), and:

$$\mathrm{Cov}\left(e_f^p, n_f^p\right) = \sigma_{\bar{e}^p,\bar{n}^p}. \tag{4.14}$$

These equations deliver $\sigma^2_{\bar{e}^p}$, $\sigma^2_{\bar{n}^p}$ and $\sigma_{\bar{e}^p,\bar{n}^p}$. Then, using these estimates, $\sigma_{\bar{q}^p,\bar{e}^p}$ and $\sigma_{\bar{q}^p,\bar{n}^p}$ are identified from:

$$\mathrm{Cov}\left(e_f^p, c_f^p\right) = \sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p} \tag{4.15}$$

$$\mathrm{Cov}\left(n_f^p, c_f^p\right) = \sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}. \tag{4.16}$$

Finally, equation (4.10) can be used to recover $\sigma^2_{\bar{q}^p}$.

Intergenerational Persistence. The intergenerational elasticity parameters $(\gamma, \theta, \rho, \lambda, \phi)$ are identified using the cross-generation moments. After using equation (4.14) to recover $\sigma_{\bar{e}^p,\bar{n}^p}$ and equation (4.8) to recover $\sigma^2_{\bar{e}^p}$, equations (4.17) and (4.20) jointly identify γ and θ. Similarly, ρ and λ are identified from (4.18) and (4.19). This leaves ϕ to be identified from equation (4.21).

$$\mathrm{Cov}\left(e_f^p, e_f^k\right) = \gamma\sigma^2_{\bar{e}^p} + \theta\sigma_{\bar{e}^p,\bar{n}^p} \tag{4.17}$$

$$\mathrm{Cov}\left(n_f^p, n_f^k\right) = \rho\sigma^2_{\bar{n}^p} + \lambda\sigma_{\bar{e}^p,\bar{n}^p} \tag{4.18}$$

$$\mathrm{Cov}\left(e_f^p, n_f^k\right) = \rho\sigma_{\bar{e}^p,\bar{n}^p} + \lambda\sigma^2_{\bar{e}^p} \tag{4.19}$$

$$\mathrm{Cov}\left(n_f^p, e_f^k\right) = \gamma\sigma_{\bar{e}^p,\bar{n}^p} + \theta\sigma^2_{\bar{n}^p} \tag{4.20}$$

$$\begin{aligned}\mathrm{Cov}\left(c_f^p, c_f^k\right) &= \phi\left(\sigma^2_{\bar{q}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{q}^p,\bar{n}^p}\right) + (\gamma+\lambda)\left(\sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) \\ &\quad + (\rho+\theta)\left(\sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right)\end{aligned} \tag{4.21}$$

Cross-sectional Parameters for Children. Finally, we employ estimates from the previous steps to identify children-specific dispersion parameters. The values of $\sigma^2_{\delta^k}$ and $\sigma^2_{\varepsilon^k}$ are identified from (4.11) and (4.12), respectively. The remaining child specific parameters are identified from the covariances of income, earnings and consumption among children (see equations 4.22 through 4.24 below) as well as from the variance of consumption in equation (4.13).

$$\mathrm{Cov}\left(e_f^k, n_f^k\right) = (\rho\gamma + \theta\lambda)\,\sigma_{\bar{e}^p,\bar{n}^p} + \gamma\lambda\sigma^2_{\bar{e}^p} + \rho\theta\sigma^2_{\bar{n}^p} + \sigma_{\delta^k,\varepsilon^k} \tag{4.22}$$

$$\begin{aligned}\mathrm{Cov}\left(e_f^k, c_f^k\right) &= \gamma(\gamma+\lambda)\,\sigma^2_{\bar{e}^p} + \theta(\theta+\rho)\,\sigma^2_{\bar{n}^p} + \phi\gamma\sigma_{\bar{q}^p,\bar{e}^p} + \phi\theta\sigma_{\bar{q}^p,\bar{n}^p} \\ &\quad + \left[\gamma(\rho+\theta) + \theta(\gamma+\lambda)\right]\sigma_{\bar{e}^p,\bar{n}^p} + \sigma^2_{\delta^k} + \sigma_{\psi^k,\delta^k} + \sigma_{\delta^k,\varepsilon^k}\end{aligned} \tag{4.23}$$

$$\begin{aligned}\mathrm{Cov}\left(n_f^k, c_f^k\right) &= \lambda(\gamma+\lambda)\,\sigma^2_{\bar{e}^p} + \rho(\theta+\rho)\,\sigma^2_{\bar{n}^p} + \phi\lambda\sigma_{\bar{q}^p,\bar{e}^p} + \phi\rho\sigma_{\bar{q}^p,\bar{n}^p} \\ &\quad + \left[\lambda(\rho+\theta) + \rho(\gamma+\lambda)\right]\sigma_{\bar{e}^p,\bar{n}^p} + \sigma^2_{\varepsilon^k} + \sigma_{\delta^k,\varepsilon^k} + \sigma_{\psi^k,\varepsilon^k}\end{aligned} \tag{4.24}$$

Over-identifying Moments. The following four inter-generational moments can be used as over-identifying restrictions for the parameter estimates:

$$\text{Cov}\left(e^p_f, c^k_f\right) = (\gamma+\lambda)\,\sigma^2_{\bar{e}^p} + \phi\sigma_{\bar{q}^p,\bar{e}^p} + (\rho+\theta)\,\sigma_{\bar{e}^p,\bar{n}^p} \tag{4.25}$$

$$\text{Cov}\left(n^p_f, c^k_f\right) = (\rho+\theta)\,\sigma^2_{\bar{n}^p} + \phi\sigma_{\bar{q}^p,\bar{n}^p} + (\gamma+\lambda)\,\sigma_{\bar{e}^p,\bar{n}^p} \tag{4.26}$$

$$\text{Cov}\left(c^p_f, e^k_f\right) = \gamma\left(\sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) + \theta\left(\sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) \tag{4.27}$$

$$\text{Cov}\left(c^p_f, n^k_f\right) = \lambda\left(\sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) + \rho\left(\sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) \tag{4.28}$$

Identification: A Graphical Example. One insight of the identification argument is that we can use elements of the covariance structure to jointly harness information about cross-sectional inequality and covariation of permanent income across generations. To illustrate how this works in practice, it helps to consider the relationships in Figure 4.1 where the y-axis measures the parental permanent earnings variance, $\sigma^2_{\bar{e}^p}$, and the x-axis represents the intergenerational earnings persistence, γ. To identify this pair of parameters we only use three empirical moments: $Var(e^p_f)$, $Cov(e^p_f, e^k_f)$ and $Var(e^k_f)$.

From moment condition (4.8), the variance of parental earnings ($\sigma^2_{\bar{e}^p}$) is uniquely identified by $\text{Var}\left(e^p_f\right)$: its value is shown as the horizontal dashed line in Figure 4.1. The moment condition (4.11) captures the tradeoff between γ and $\sigma^2_{\bar{e}^p}$, holding constant other persistence and variance parameters (i.e., $\theta, \sigma^2_{\bar{n}^p}, \sigma_{\bar{e}^p,\bar{n}^p}$ and $\sigma^2_{\delta^k}$). This is plotted as the negatively sloped dotted line in Figure 4.1. The intersection of the dotted line with the dashed line uniquely identifies the persistence parameter, γ. However, our model features an additional restriction: the exact location of the pair $\left(\gamma, \sigma^2_{\bar{e}^p}\right)$ needs to be consistent with the moment condition (4.17), imposing an additional tradeoff between the two parameters (shown by the solid line). That is, $\sigma^2_{\bar{e}^p}$ and γ must be such that both the solid and the dotted lines intersect the dashed line at a common location. One can verify that the location where all three moment conditions hold in Figure 4.1 corresponds to the baseline parameter estimates presented in Section 4.4.

4.3.2 Estimation

We estimate model parameters using a generalized method of moments that minimizes the sum of squared deviations between empirical and theoretical second moments. We use an equally weighted distance metric because of the small sample biases associated with using a full variance-covariance matrix featuring higher-order moments (see Altonji and Segal, 1996). Data on earnings, other income and consumption is used to calculate the empirical moments, after removing time and birth-cohort effects.[66]

In Appendix C.5.1, we decompose the total variation in the baseline data into a component explained by observable characteristics, like race, education, number of family members, employment

[66]Empirical moments are constructed using the residuals of a log-linear regression of the variables on a full set of year and cohort dummies. This is done separately for the parent and child generations.

status, etc., and a residual component representing heterogeneity in unobservable factors. This is helpful to establish whether the persistence and transmission of inequality across generations are due to observable or unobservable characteristics of the parent-child pairs.

4.3.3 Data

We use data from the Panel Study of Income Dynamics (PSID). This dataset is widely used in the analysis of intergenerational persistence of economic outcomes because the offspring of original sample members become part of the survey sample when they establish independent households. Using these data has also the advantage of making our analysis easily comparable to existing studies based on the PSID. We focus on the nationally representative sample of the PSID (from the Survey Research Centre, SRC) between 1967 and 2014, and exclude samples from the Survey of Economic Opportunity (SEO), immigrant and Latino sub-populations. To avoid noise due to weak labour market attachment and variation in marital status, we sample married households with a male head and at least 5 years of data.[67] We also restrict the sample to families with non-negative labour earnings and total income, that do no more than 5,840 hours of work in any year, and with wages at least half of the federal minimum wage. Finally, we select out households that experience annual earnings growth of more than 400%. Baseline results focus on intergenerational linkages between fathers and sons.[68] For each generation, we consider income and expenditure from age 25 onwards, with a maximum sample age of 65, to avoid issues related to retirement choices. By design, the income and consumption information of parents refers to later stages of the life cycle. In our baseline sample of 760 unique father-son pairs, the average parental age is 47 years while that for the children is only 37 years. Further details about data and sampling restrictions are in Appendix C.2.

Labour earnings data for the household head and his spouse are readily available for all survey waves of the PSID. Data about transfers from public and private sources for the husband and the wife are also available for most years. In contrast, the consumption expenditure data can be sparse, and not presented as a single variable in the PSID. Expenditures on food are the only category that is observed almost consistently since the earliest 1968 wave, and we use food outlays as the consumption measure for the baseline estimation. In Section 4.6, we examine the robustness of our findings to an alternative consumption measure, suggested by Attanasio and Pistaferri (2014), that relies on 11 major categories of consumption outlays that are reported from 1999 onwards. This approach measures total consumption expenditure at the household level by estimating a simple demand system using data for the years in which all 11 consumption expenditures were available in the PSID; then, by inverting the demand system, one can recover total expenditures for the years before 1999. The method relies on the theory of consumer demand and two-stage budgeting: the allocation of resources spent in a given period over different commodities is assumed to depend

[67] The restriction to married households is helpful but not inconsequential, as intergenerational insurance may come into play exactly at the time of relationship breakdown (Fisher and Low, 2015).

[68] Our focus on father-son linkages also avoids the sample issues discussed in Hryshko and Manovskii (2019).

on relative prices, taste-shifters (demographic and socio-economic variables) and total expenditure. Details about the variables, their availability in the survey and the demand system estimation procedure are reported in Appendix C.2. We adjust data on household-level expenditures through the OECD adult equivalence scale.

4.4 Results

Table 4.1 reports the variances of earnings, other income and consumption expenditures for parents and children.[69] These variances, along with the empirical moments reported in Appendix Figure C.2, are used in the baseline implementation to estimate intergenerational persistence parameters and the underlying variance-covariance structure of permanent income and consumption for each generation. We summarize the within-sample fit of the model in Figure C.2 of Appendix C.4.

The two lifetime income sources are much more dispersed than expenditures for both the generations, indicating the presence of mechanisms that induce cross-sectional consumption smoothing. This may occur through both formal taxes and transfers and the heterogeneous saving and spending behaviour of households. Amongst income sources, labour earnings of the male household head are less dispersed than other family income, which consists of transfer income and wife earnings. In Appendix Table C.14 we show that the higher dispersion of other income is due to the uneven distribution of transfer income; in contrast, spousal earnings are significantly less dispersed than transfers in both generations. The relative magnitudes of head earnings and family consumption dispersion reported in Table 4.1 are similar to those found in studies by Krueger and Perri (2006) and Attanasio and Pistaferri (2014). For a direct comparison with these studies that do not split the data into two generations, in Appendix Figure C.1 we show the evolution of cross-sectional earnings and consumption inequality for the last four decades in the U.S.

The age range used to calculate these variances is wider for parents than it is for children since parents are observed for a longer period in PSID data. Therefore, differences in the magnitude of variances of parents and children, shown in Table 4.1, do not imply a decline in income inequality across generations. Rather, these differences reflect shocks accruing at different stages of the life-cycle. Table 4.7 in Section 4.5 reports variances based on samples where the ages of both parents and children are restricted between 30 and 40. These variances illustrate the evolution of inequality across generations, showing a relative increase in inequality among children that is consistent with the well-established notion of increasing income U.S. inequality over the past decades. The age restriction, however, substantially reduces the sample size, and in the baseline analysis, we use the wider age range for parents to obtain more accurate estimates of parental permanent income. Since we do not observe children in the later part of their working lives, our estimates reflect how parental heterogeneity impacts dispersion among children in the earlier decades of their adult lives.

[69] In Appendix Table C.3, we decompose each statistic into the variance due to observable characteristics of a generation and the residual variance due to unobservable factors.

Intergenerational Elasticities. Table 4.2 reports estimates of intergenerational persistence parameters. The elasticity is highest for earnings, with the pass-through γ estimated at 0.23; in contrast, the elasticity for other household income, ρ, is 0.10 and that for consumption, ϕ, is 0.15. It is important to emphasize that the significant covariation in idiosyncratic expenditure shifters q across generations, captured by the parameter ϕ, contributes to consumption inequality over and above any effects working through the earnings and other income channels. That is, family influences on consumption expenditures build-up through three inter-dependent channels: earnings, other household income, and persistence in consumption and saving propensities.

Higher parental earnings are associated with higher levels of other income among offspring, with the cross-elasticity λ equal to 0.21: this positive covariation holds for both transfers and spousal earnings among children (see Table 4.5). On the other hand, other household income has little effect on children's earnings, with the elasticity θ estimated to be small albeit statistically significant. Explicitly accounting for these cross-effects between different dimensions of intergenerational pass-through (namely, male head earnings, wife earnings and transfer income) turns out to be an important contribution of our approach over the standard reduced-form analysis of intergenerational persistence. As we show in Section 4.6, ignoring these cross-effects may lead to misleading inference about the role of family influences for cross-sectional inequality in the children's generation.

In Appendix C.5.1, we show that all pass-through parameters in Table 4.2 are primarily driven by intergenerational persistence in observable characteristics like race, employment status, educational attainment, family structure, state of residence and so on. In particular, we document that education accounts for a large component of the earnings pass-through across generations, corroborating the evidence on this transmission channel in Landersø and Heckman (2017).[70]

Permanent Income and Consumption. Table 4.3 reports estimates of the variances and covariances of the permanent components of earnings, other income and consumption shifters.[71] The importance of jointly estimating income and consumption processes becomes apparent when examining these estimates. To illustrate how covariations are key to account for data patterns, we note that the variance of the permanent consumption components, $\sigma_{\bar{q}}$ is larger than that of permanent earnings in both generations; however, we know that consumption expenditures are much less dispersed than earnings. This apparent discrepancy highlights the role of the negative covariation between permanent earnings and idiosyncratic consumption shifters. Estimates of this covariance are -0.27 for the parents' generation (see $\sigma_{\bar{e},\bar{q}}$) and exhibit a similar magnitude in the child's generation. In addition, the permanent component of other income exhibits even stronger negative covariation with lifetime consumption shifters (see $\sigma_{\bar{e},\bar{q}}$). The negative covariation between the permanent components of consumption and income mitigates the impact of income inequality on consumption inequality; that is, the negative covariances compress the distribution of log consump-

[70] Empirical support for this result can also be found in the mobility matrix of education in Appendix Table C.8. High persistence of education across generations might be due to unobserved ability, or to higher human capital investments by parents with higher earnings (see Lefgren, Sims, and Lindquist, 2012).

[71] In Appendix Table C.5 we report estimates of these parameters separately for cases where we use variation in the outcome variables that are either explained by observable characteristics or left unexplained by them.

tion and drive its overall variance below the variance of income. Moreover, these estimates suggest that higher-income families save proportionally more and have, on average, a lower propensity to consume.[72] Such traits are passed across generations, which reinforces their mitigating influence on consumption dispersion.

4.4.1 Role of Parental Heterogeneity

The quantitative importance of parental heterogeneity for offspring depends on three aspects: (i) intergenerational persistence, (ii) the level of inequality in the parents' generation, and (iii) the magnitude of idiosyncratic heterogeneity among kids. We gauge the influence of parental factors in two ways: first, we compute the share of earnings, income and consumption variances that is explained by pre-determined parental heterogeneity; second, we show how the cross-sectional distributions of these outcomes change if differences in parental characteristics are removed.

Variance Accounting. Table 4.4 summarizes the impact of parental heterogeneity on the variance of children outcomes. Let $\text{Var}[y^k(p)]$ measure the offspring variance that is explained by parental factors for variable $y \in \{e, n, c\}$, while $\text{Var}[y^k]$ denotes the total cross-sectional variance in the kids' generation. The ratio $\frac{\text{Var}[y^k(p)]}{\text{Var}[y^k]}$ quantifies the share of total variation attributed to parental heterogeneity.[73]

Combining the estimates in Tables 4.2 and 4.3, we are able to break down the relative contributions of parental and idiosyncratic heterogeneity to the cross-sectional dispersion of child outcomes (see Table 4.4). By far the largest impact of parental heterogeneity is on consumption dispersion, as it accounts for almost 30% of total variation among offspring. Parental factors account for much less of the variation in income — 8% and 4% for earnings and other income, respectively. As discussed before, this is consistent with the observation that intergenerational transmission of consumption and saving behaviours, after accounting for the level of income, is an important channel of intra-family persistence in consumption expenditures. Since the cross-sectional distribution of expenditures is more compressed than its counterparts for earnings and other income, parental influences end up explaining a much larger share of this lower variance. Nevertheless, it is clear that idiosyncratic heterogeneity accruing over the life cycle accounts for most of the dispersion of income and consumption outcomes in the younger generation.[74]

Lastly, it is important to emphasize that a significant share of parental influence on consumption dispersion can only be identified if one allows for non-zero cross-elasticities λ and θ between earnings and other income in the two generations. Restricting these cross-elasticities to zero not

[72] See Abbott and Gallipoli (2019) and Straub (2018) for recent evidence of high saving rates among the rich. Fan (2006) suggests that this maybe motivated by bequest motives. Dynan, Skinner, and Zeldes (2004) argue that other non-bequest motives account for this excess savings.

[73] For example, the contribution of parents in the cross-sectional earnings variance in the kids' generation is given by the ratio $\frac{\gamma^2\sigma^2_{\bar{e}^p}+\theta^2\sigma^2_{\bar{n}^p}+2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p}}{\sigma^2_{\delta^k}+\gamma^2\sigma^2_{\bar{e}^p}+\theta^2\sigma^2_{\bar{n}^p}+2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p}}$. For an illustration of all the calculations involved, see Appendix C.5.2.

[74] In Appendix Table C.6, we document that most of the explanatory power of parental heterogeneity is due to observable characteristics. In contrast, parental factors appear at most marginal in explaining unobservable heterogeneity in the younger generation.

only diminishes the quantitative contribution of parental heterogeneity to consumption dispersion but also artificially boosts the parental importance for earnings heterogeneity. This highlights again the co-dependence of these processes and the biases that are introduced by ignoring it. The mechanism behind these biases is discussed in Section 4.6 where we re-estimate the model after restricting $\lambda = \theta = 0$.

Marital Selection. In the baseline model, other income is the sum of transfer income (both public and private transfers) and spousal earnings. Table 4.5 reports estimates of intergenerational pass-through elasticities under the restriction that other income consists only of transfers (column 1) or spousal earnings (column 2).[75] Focusing, in turn, on transfer or spousal income alone alters the sample size because of missing values. Therefore, to aid comparison, we re-estimate the baseline model (with the broader measure of other income) on the sub-sample for which we have non-missing observations for both transfers and wife's earnings. We report the latter estimates in column (3).

When only spousal earnings are used in estimation, all intergenerational elasticity estimates are strongly significant and at least as large as their baseline counterparts. The point estimates of intergenerational persistence for consumption shifters and other income are roughly 50% higher than baseline results. In contrast, persistence parameters are low and very imprecisely estimated in the specification featuring transfers alone. By removing transfer income, which is rather noisy, we effectively focus on the most significant and better-measured component of other household income.

The fact that spousal earnings are rather persistent across parent-child pairs (with a ρ elasticity of 0.14) suggests that spousal sorting may partly depend on family traits and that maternal earnings may play a role in female partner choices, especially for what pertains to wife's labour market participation and earnings. This is consistent with findings in Fernandez, Fogli, and Olivetti (2004), who document preference formation based on maternal characteristics.

In Table 4.6 we break down children inequality into parental and idiosyncratic components for specifications featuring spousal income only (Panel A) and transfer income only (Panel B), and compare these to what we obtain for the baseline model estimated on the same sample (Panel C). It is instructive to notice that removing transfer measures marginally increases estimates of family persistence on earnings of both spouses. This indicates that transfers somewhat confound and offset parental influences on total household earnings. In part, these larger estimates are due to the significantly lower variance of other income when noisy transfer measures are omitted. In contrast, the contribution of family background to consumption dispersion is almost the same as in the baseline, suggesting that earnings and consumption elasticities are generally sufficient to characterize the dispersion of children expenditures.

Even after shutting down direct family transfers, parental influences on consumption inequality are estimated to be stronger than on household income, albeit the gap is smaller than in the benchmark model. This suggests that family influences on expenditure operate through the trans-

[75] In Table C.13 of the appendix we report the associated variance-covariance estimates.

mission of earning potential, rather than through direct transfers. Moreover, earning potential is increased at the household level, as parents affect both the offspring's earning ability and their marital choices. It appears that marital selection, together with the ability to earn and the persistence of consumption and saving behaviours, can account for most of the parental influences identified in the baseline.

4.4.2 Counterfactual Cross-sectional Distributions

Absence of intergenerational transmission is equivalent to a setting with randomly matched parent-child pairs. A simple way of gauging the impact of family background in this setting is to plot the observed and counterfactual cross-sectional distribution of each outcome in the children generation (top panels of Figure 4.2) and their local differences, (as measured by the histograms in the bottom panels of Figure 4.2). The histograms represent, for each interval of the domain, the probability mass of the actual distribution minus the corresponding mass in the counterfactual.[76]

The counterfactual distributions are visibly less dispersed, with the strongest departure from baseline observed for lifetime consumption. In all counterfactuals, much of the probability mass at the lower tail is reshuffled towards the middle of the distribution, while the right tail also becomes somewhat less heavy. This suggests that the impact of parental heterogeneity is especially relevant at the bottom of the distributions, where it stretches the left tail of both income and consumption. These findings are in line with results in Table 4.4 and provide additional evidence that family influences account for a comparatively larger share of consumption variation.

4.5 The Evolution of Inequality across Generations

The magnitudes of the intergenerational pass-through parameters and the idiosyncratic variances raise questions about the evolution of inequality across generations. A longer data panel would be ideal to identify persistence across multiple generations, since the current span of PSID data covers, at most, the working life of children born between the 1950s and the early 1980s. This makes it hard to obtain direct estimates of the impact of grandparents on grandchildren and generations further apart. However, under a stationarity assumption, one can examine the projected path of inequality by computing a first-order approximation of the expected evolution of the variances of income and consumption starting from current levels.[77]

To examine inequality across generations, we compute the long-run steady-state variances of earnings, other income and consumption. These measures describe how dispersed income and consumption would be if, all else equal, the baseline model were allowed sufficient time to converge to its steady state. By comparing current variances to their steady-state values, one can tie changes in cross-sectional inequality to the intergenerational persistence of parental advantage. Since it is

[76] Appendix C.5.2 describes how we measure the actual and counterfactual distributions. In the counterfactual case, all parental channels are shut down. All exercises rely on estimates from Tables 4.2 and 4.3.

[77] Stationarity implies holding $\{\sigma^2_{\delta^k}, \sigma^2_{\psi^k}, \sigma^2_{\varepsilon^k}, \sigma_{\delta^k,\varepsilon^k}, \sigma_{\psi^k,\varepsilon^k}, \sigma_{\psi^k,\delta^k}\}$ constant at their baseline values. Of course, large changes in structural parameters might mitigate or exacerbate the baseline scenario.

preferable to focus on individuals of similar age across generations and we do not observe children in the second half of their working lives, we restrict the age of both parents and children in our sample to be between 30 and 40.[78]

A Vector Representation of the Model. Earnings, other income and consumption shifters evolve through generations of family f according to the following vector autoregressive process:

$$\begin{bmatrix} \bar{e}_f^{k_t} \\ \bar{n}_f^{k_t} \\ \bar{q}_f^{k_t} \end{bmatrix} = \begin{bmatrix} \gamma & \theta & 0 \\ \lambda & \rho & 0 \\ 0 & 0 & \phi \end{bmatrix} \cdot \begin{bmatrix} \bar{e}_f^{k_{t-1}} \\ \bar{n}_f^{k_{t-1}} \\ \bar{q}_f^{k_{t-1}} \end{bmatrix} + \begin{bmatrix} \delta_f^{k_t} \\ \varepsilon_f^{k_t} \\ \psi_f^{k_t} \end{bmatrix}.$$

The superscript $\{k_t\}$ identifies the t^{th} generation of kids. Since k_1 denotes the first generation of kids, we define k_0 to be the parents' generation in our data, that is, $\bar{x}_f^{k_0} \equiv \bar{x}_f^p$ for any variable $x \in \{e, n, q\}$. The joint distribution of the covariance-stationary idiosyncratic shocks is

$$\begin{bmatrix} \delta_f^{k_t} \\ \varepsilon_f^{k_t} \\ \psi_f^{k_t} \end{bmatrix} \sim N \left[\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2_{\delta^k} & \sigma_{\delta^k,\varepsilon^k} & \sigma_{\delta^k,\psi^k} \\ \sigma_{\delta^k,\varepsilon^k} & \sigma^2_{\varepsilon^k} & \sigma_{\varepsilon^k,\psi^k} \\ \sigma_{\delta^k,\psi^k} & \sigma_{\varepsilon^k,\psi^k} & \sigma^2_{\psi^k}. \end{pmatrix} \right]$$

Using parameter estimates, we simulate the VAR forward, iterating until convergence.[79] This delivers simulated data series for $\bar{e}_f^{k_t}$, $\bar{n}_f^{k_t}$, $\bar{q}_f^{k_t}$, $\delta_f^{k_t}$, $\varepsilon_f^{k_t}$ and $\psi_f^{k_t}$. To obtain a series for log consumption, we use the relationship:

$$c_f^{k_t} = \phi q_f^{k_{t-1}} + (\gamma + \lambda) e_f^{k_{t-1}} + (\rho + \theta) n_f^{k_{t-1}} + \delta_f^{k_t} + \varepsilon_f^{k_t} + \psi_f^{k_t},$$

for $t \geq 1$. Having recovered the (log) series for the permanent components of earnings, other income, and consumption, we calculate their long-run variances and report them in column (3) of Table 4.7.

Current versus Long-Run Inequality. Comparing steady-state variances with those observed in data, we see that for earnings and consumption the inequality in the parents' generation is the lowest (see column 1 of Table 4.7), followed by that in the children's generation (column 2 of Table 4.7). Steady-state inequality is the largest, suggesting that the variance of lifetime earnings and consumption expenditures might rise further from current levels. For other income, inequality in the children's generation is lower than in their parents' generation and slightly lower than its value in the steady-state. The steady-state variances implied by the baseline model are not far above what is measured in the children's generation. This observation reflects the low value of the estimated pass-through parameter γ, meaning that predicted long-run inequality reflects primarily the variance of idiosyncratic shocks.

[78] Baseline estimates are based on a larger sample that includes observations for older parents.

[79] Since we restrict the age range between 30 and 40 years, we re-estimate the baseline model on a smaller sample. The estimates are reported in column (1) of Tables 4.10 and C.12. The VAR is simulated for 100,000 generations.

To illustrate the quantitative importance of intergenerational elasticities in the long-run, we re-estimate the baseline model using a constrained version of the GMM estimator where we hold constant the earnings persistence γ at pre-determined values. By exogenously setting larger or smaller values of γ, we can assess whether, and how much, steady-state inequality might deviate from its initial value. Table 4.8 shows that for counterfactually high values of γ, earnings inequality in the children generation (column 4) can be substantially different from long-run model outcomes (column 5). Moreover, a trade-off between inter-generational persistence, γ (column 1) and idiosyncratic heterogeneity, $\sigma^2_{\delta^k}$ (column 2) is evident when explaining the total child variance (column 4).[80]

Despite a falling variance for idiosyncratic innovations, $\sigma^2_{\delta^k}$, steady-state inequality in column 5 increases with the magnitude of γ. Thus, the cross-generational persistence, rather than the innovations variance, emerges as the key determinant of long-run inequality and as the main reason for the similarity of $\text{Var}\left(e^k\right)$ and $\text{Var}\left(e^*\right)$.[81]

These results emphasize that, without any increases in the underlying dispersion of idiosyncratic innovations, one would have to assume implausibly large values of the intergenerational pass-through to induce significantly higher long-run inequality. It follows that intergenerational persistence dictates the proportional impact of parental heterogeneity on inequality. Further evidence of this is in the last column of Table 4.8, which documents how changes in γ lead to significant variation in the contribution of parental factors to cross-sectional earnings inequality. A larger γ amplifies the contribution of family background: the parental contribution to inequality swings widely, between 1% and 12% (for values of γ between 0.1 and 0.4) even when steady-state earnings dispersion $\widehat{Var(e^*)}$ barely changes.

It is interesting to contrast the values in column 6 of Table 4.8 with baseline estimates of the importance of parental factors in Table C.6, where the age range was not restricted. Restricting the age range over which parents' income is measured implies that the importance of family background declines from about 8% to 4% of total variation: that is, roughly half of the parental impact on inequality among children accrues by the time parents reach age 40.

A final caveat for these results is that inference about the evolution of inequality is based on stationary parameter estimates. For this reason in Appendix C.6 we consider the implications of changes in structural parameter estimates on inequality going forward and we explore how inequality evolves over subsequent generations (parent, child, grandchild) while converging to its steady-state

[80] When intergenerational persistence γ is set to a higher value, the GMM estimator delivers a lower variance of idiosyncratic heterogeneity (e.g., for earnings, lower $\sigma^2_{\delta^k}$) since the amount of cross-sectional inequality among children is fixed by what is observed in the data.

[81] A striking feature of the GMM estimates in Table 4.8 is that the child variance remains constant and matches exactly the empirically observed variance. On the other hand, the empirically observed parental variance is 0.199 and none of the estimates matches this figure exactly. To understand this, consider the relevant moment conditions. At first glance, equation (4.8) suggests that the parental variance estimate should be independent of the choice of γ. However, the GMM tries to satisfy equation (4.17), which implies a direct trade-off between γ and $Var(e^p)$. Thus, increasing γ tends to decrease $Var(e^p)$. On the other hand, whatever the values for γ and $Var(e^p)$, the empirically observed value of $Var(e^k)$ can always be matched exactly by choosing the free parameter $\sigma^2_{\delta^k}$, which does not enter any other moment condition.

level.

4.6 Robustness and Extensions

To assess the robustness of these findings we perform three types of sensitivity checks. We begin by examining whether a specific cohort drives the baseline results. In the second round of checks, we estimate variants of the baseline model where we (i) restrict specific channels of intergenerational transmission, (ii) consider a sample of randomly matched parent-child pairs, (iii) employ alternative measures of expenditure, and (iv) target additional moments in the GMM estimation. Third, we estimate an alternative model of intergenerational persistence, specified in terms of growth rates of the outcome variables. This allows one to draw inference about intergenerational persistence of idiosyncratic innovations to income and consumption expenditures.

4.6.1 Estimates by Child Birth-Cohort

As a first check, we split observations by child birth-cohort. Grawe (2006) and Gouskova, Chiteji, and Stafford (2010) argue that life-cycle bias may be important when estimating intergenerational persistence. Since children from different birth cohorts are observed for varying lengths of their life-cycle in the PSID, we avoid life-cycle bias by restricting the age of parents and kids to be between 30 and 40 years.

Table 4.9 shows the cross-sectional variances of economic outcomes for the parents and kids for different child birth-cohorts.[82] Contrary to the estimates in Table 4.1 where we imposed no age restrictions, we find that controlling for the life-cycle bias through age restriction makes the dispersion of earnings higher in the child generation for all cohorts, consistent with the empirical literature documenting growing inequality in the U.S. over our sample period.

Table 4.10 presents estimates of intergenerational pass-through parameters by children's decade of birth. The results are qualitatively similar to the baseline ones.[83] The differences between estimates of the intergenerational pass-through parameters for the 1960s and 1970s cohorts are not statistically significant. In Appendix Table C.11 we consider whether the importance of parental influence in explaining cross-sectional heterogeneity in the child generation varies by children cohorts. Contrasting the 1960s and 1970s cohorts, the contribution of parental heterogeneity changed only for consumption, dropping from about 38% to 16%. However, cohort-specific sample sizes are small enough to suggest caution when comparing these shares.

4.6.2 Additional Robustness Checks

Next, we perform four additional robustness checks. First, we estimate a restricted version of the baseline model where we shut down cross-effects in intergenerational persistence; that is, we set λ

[82] We do not report estimates for the 1950s cohort because the sample size for that cohort becomes too small and estimates are very noisy.

[83] Some of the parameter estimates lose statistical significance, as the age restrictions result in a much smaller sample, weakening the precision of the estimates.

and θ to zero. Second, we perform a placebo test where we randomly match parents and children and show that baseline estimates capture genuine intergenerational linkages rather than spurious correlations. Third, we use imputed consumption expenditures, rather than food expenditure, to measure household outlays. Finally, we re-estimate the model using both cross-sectional and panel variation: this extension increases the number of moments, as well as the number of parameters to estimate. The main results from these robustness checks are presented in Table 4.11, with additional details in Appendix Table 4.12.

Restricting Cross-Effects between Income Sources

We consider a restricted version of the baseline model that does not allow for any effect from parental earnings on other income of the children, nor from parent's other income onto the child's earnings, that is, imposing both λ and θ to be zero. Under these restrictions the point estimates of the parameters change significantly, overstating the importance of parents for earnings inequality among children. Most of the difference from the baseline estimates can be attributed to the restriction that $\lambda = 0$, as in the baseline model θ is already close to zero. By restricting λ to be zero, one effectively decreases its value below the positive baseline estimate. This mechanically pushes up the estimate of γ to guarantee a fairly constant value of $(\gamma + \lambda)$, the total intergenerational persistence from parental earnings to child outcomes. Hence, the exercise highlights the importance of allowing for cross-effects above and beyond the direct channels captured by γ and ρ when drawing inference about pass-through parameters. Table 4.13 also illustrates that parental heterogeneity explains much less of cross-sectional consumption dispersion when λ and θ are set to zero; this confirms that, while higher parental earnings have a positive direct impact on the earnings and expenditures of children, the consumption distribution is shaped by several other forces and ignoring the indirect effects among different sources of income can lead to incorrect inference.

Placebo Test: Random Matching of Parents and Children

One might worry that spurious correlations in the data affect estimates of parent-child pass-through parameters. To account for this possibility one can perform a placebo test using a sample in which parents and children are randomly matched. Estimates based on this sample imply no role of parental heterogeneity for inequality in the children's generation, as seen in column (3) of Table 4.13. The lack of significance in the randomly matched sample indicates that genuine family linkages, rather than spurious correlations, drive baseline estimates.

Alternative Measures of Expenditure

The baseline analysis uses food expenditures as the consumption measure. This choice is dictated by necessity as food consumption is available throughout the PSID sample. However, alternative components of consumption might exhibit different intergenerational persistence. We examine the importance of other expenditure categories in two ways. First, we impute total consumption using the procedure suggested by Attanasio and Pistaferri (2014); this approach exploits rich consumption expenditure information available in the PSID after 1997 to approximate households' outlays in the earlier years of the survey. We report results for this alternative consumption measure in column (4) of Tables 4.11, 4.13 and 4.12. Estimates based on this broader gauge of expenditures suggest a stronger role of parental heterogeneity for consumption dispersion among children, with roughly half of the total dispersion due to family linkages.[84]

In a second sensitivity exercise, we restrict the sample to the post-1997 period, when there is no need for imputation of non-food consumption. Estimates from this smaller sample suggest a parental contribution to consumption inequality of roughly 24%, comparable to the baseline estimate based on food consumption alone.[85] Inevitably, the smaller size of the post-1997 sample makes estimates less precise.

Additional Moments and Parameters: Panel Variation

In the baseline analysis, we average across yearly observations for each sample member and do not account for year-to-year individual variation. Time-averaging significantly reduces the impact of classical measurement error, but it also precludes identification of the variances of mean-zero transitory shocks to earnings, income and consumption.[86] Thus, accounting for panel variation introduces extra parameters due to the need to estimate the variances and covariances of per-period transitory shocks. By the same token, this extra information introduces new moment restrictions. In Appendix C.7.3 we report the full set of moments and parameters. As shown in column (5) of Tables 4.11, 4.13 and 4.12, modelling period-specific variation makes little difference. However, standard errors are visibly inflated, as one would expect when measurement error becomes more severe.

4.6.3 An Alternative Model of Intergenerational Persistence

In the baseline model specification, cross-sectional heterogeneity is partly inherited through intergenerational persistence of individual fixed effects in income and consumption. An alternative hypothesis is that these linkages may occur through persistence in growth rates. To examine this possibility, we examine a model in which the permanent components of both earnings and other income are random walk processes (see for example, Blundell, Pistaferri, and Preston, 2008), and the contemporaneous permanent innovations to these processes are correlated across parent-child pairs. Appendix C.8 presents details of this alternative model, along with the identification strategy and parameter estimates. We find little or no evidence of intergenerational persistence in permanent innovations to earnings, other income and consumption expenditures.

[84] Arguably, this higher estimate of the parental contribution to consumption inequality is an upper bound of the true value, as it may partly reflect latent persistence of observable characteristics used to impute consumption.

[85] Estimates are not reported in the table and are available upon request.

[86] When using panel variation, classical measurement error is indistinguishable from per-period transitory innovations. Estimates of variances and covariances of transitory shocks are presented in Table C.15.

4.7 Conclusion

This chapter examines the importance of heterogeneity among parents for understanding income and consumption inequality. We estimate the intergenerational elasticities of earnings, other income and consumption and document their significance for the persistence of inequality across generations. We find that the quantitative contribution of idiosyncratic life cycle shocks to inequality is much larger than the contribution of parental effects. This implies that families provide limited insurance against idiosyncratic life-cycle risk.

In reaching this conclusion, we highlight the importance of jointly estimating the income and expenditure processes, and of accounting for cross-effects between different sources of income and consumption. As an example of these cross-effects, we document a negative covariation between permanent income and consumption shifters, which suggests that higher-income families have a lower average propensity to consume.

Our estimates imply that intergenerational persistence is not by itself high enough to induce further large increases in inequality over time and across generations. This reiterates the prominent role of idiosyncratic life-cycle risk, which diffuses and attenuates the impact of family background on the cross-sectional distributions of life-cycle income and consumption.

4.8 Tables

Table 4.1: Variances

Variables	**Parent**	**Child**
Earnings	0.291	0.248
Other Income	0.808	0.534
Consumption	0.097	0.114

Table 4.2: Estimates: Intergenerational Elasticities

Variables	**Parameters**	**Estimates** (1)
Earnings	γ	0.230
		(0.027)
Other Income	ρ	0.100
		(0.023)
$\bar{e}^p_f$ on $n^k_{f,t}$	λ	0.206
		(0.032)
$\bar{n}^p_f$ on $e^k_{f,t}$	θ	0.055
		(0.019)
Consumption Shifters	ϕ	0.154
		(0.032)
No. of Parent-Child Pairs	N	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. Data is purged of year and birth-cohort effects. The average age for parents in the sample is 47 years; that of children is 37 years.

Table 4.3: Estimates: Variances and Covariances of Idiosyncratic Components

	Parameters	**Estimates**
		(1)
Parental Outcomes: Variances		
Permanent Earnings	$\sigma^2_{\bar{e}^p}$	0.295
		(0.021)
Permanent Other Income	$\sigma^2_{\bar{n}^p}$	0.806
		(0.06)
Permanent Consumption Shifters	$\sigma^2_{\bar{q}^p}$	1.031
		(0.065)
Child Idiosyncratic Heterogeneity: Variances		
Permanent Earnings	$\sigma^2_{\delta^k}$	0.228
		(0.011)
Permanent Other Income	$\sigma^2_{\varepsilon^k}$	0.511
		(0.043)
Permanent Consumption Shifters	$\sigma^2_{\psi^k}$	0.730
		(0.056)
Parental Outcomes: Covariances		
Consumption Shifters & Earnings	$\sigma_{\bar{q}^p,\bar{e}^p}$	-0.271
		(0.024)
Consumption Shifters & Other Income	$\sigma_{\bar{q}^p,\bar{n}^p}$	-0.818
		(0.061)
Earnings and Other Income	$\sigma_{\bar{e}^p,\bar{n}^p}$	0.070
		(0.013)
Child Idiosyncratic Heterogeneity: Covariances		
Consumption Shifters & Earnings	σ_{ψ^k,δ^k}	-0.247
		(0.018)
Consumption Shifters & Other Income	$\sigma_{\psi^k,\varepsilon^k}$	-0.522
		(0.048)
Earnings & Other Income	$\sigma_{\delta^k,\varepsilon^k}$	0.075
		(0.013)
No. of Parent-Child Pairs	N	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. Data is purged of year and birth-cohort effects. The average age for parents in the sample is 47 years; that of children is 37 years.

Table 4.4: Breaking Up Child Inequality: Parental versus Idiosyncratic Heterogeneity

Variables	**Child Variance**	**Variance due to Parents**	**Idiosyncratic Variance**
	(1)	(2)	(3)
Earnings	0.248	0.020 (8.1%)	0.228 (91.9%)
Other Income	0.534	0.024 (4.4%)	0.510 (95.6%)
Consumption	0.114	0.034 (29.8%)	0.080 (70.2%)

Note: Numbers obtained using parameter estimates from Tables 4.2 and 4.3.

Table 4.5: Decomposition of Other Income: Intergenerational Elasticity Estimates

	Parameters	**Just Transfers**	**Spouse Earnings**	**Other Income**
		(1)	(2)	(3)
Earnings	γ	0.239	0.275	0.254
		(0.050)	(0.027)	(0.032)
Other Income	ρ	0.031	0.142	0.097
		(0.046)	(0.036)	(0.033)
$\bar{e}^p_f$ on $n^k_{f,t}$	λ	0.107	0.232	0.184
		(0.073)	(0.033)	(0.045)
$\bar{n}^p_f$ on $e^k_{f,t}$	θ	-0.007	0.144	0.086
		(0.017)	(0.033)	(0.027)
Consumption Shifters	ϕ	0.007	0.372	0.217
		(0.047)	(0.047)	(0.047)
No. of Parent-Child Pairs	N	459	459	459

Note: Bootstrap standard errors (100 repetitions) reported in parentheses. Food expenditures used as a measure of consumption. Appendix Tables C.13 and C.14 present estimates of the corresponding variance-covariance parameters.

Table 4.6: Parental versus Idiosyncratic Heterogeneity: Role of Marital Selection

Variable	**Child Variance** (1)	**Variance due to Parents** (2)	**Idiosyncratic Variance** (3)
Panel A			
Earnings	0.229	0.033 (14.6%)	0.196 (85.4%)
Wife Earnings	0.322	0.026 (8.1%)	0.296 (91.9%)
Consumption	0.113	0.025 (22.5%)	0.088 (77.5%)
Panel B			
Earnings	0.229	0.016 (7.1%)	0.213 (92.9%)
Transfer Income	1.068	0.005 (0.4%)	1.063 (99.6%)
Consumption	0.113	0.034 (30.3%)	0.079 (69.7%)
Panel C			
Earnings	0.229	0.024 (10.7%)	0.205 (89.3%)
Other Income	0.457	0.016 (3.5%)	0.441 (96.5%)
Consumption	0.113	0.027 (24.2%)	0.086 (75.8%)

Note: Panel A corresponds to the case where wife earnings is used as the measure of other income. Panel B uses transfer income of the male household head and his wife as the measure of other income. Panel C corresponds to the case where other income is defined as the sum of wife earnings and total transfer income. Numbers obtained using parameter estimates in column (2) of Tables 4.5 and C.13, based on sample of 459 unique parent-child pairs in which both transfers and wife earnings are not missing.

Table 4.7: Steady-State versus Current Inequality

Variables	**Parental Inequality** (1)	**Child Inequality** (2)	**Steady-state Inequality** (3)
Earnings	0.199	0.251	0.255
Other Income	0.845	0.669	0.676
Consumption	0.097	0.118	0.127

Note: Estimates based on sample of 336 unique parent-child pairs. Age restricted between 30 and 40 years.

Table 4.8: Importance of Parents: Varying Persistence γ

γ	$\widehat{\sigma^2_{\delta^k}}$	$\widehat{Var(e^p)}$	$\widehat{Var(e^k)}$	$\widehat{Var(e^*)}$	$\frac{\gamma^2 \widehat{Var(e^p)}}{\widehat{Var(e^k)}}$
(1)	(2)	(3)	(4)	(5)	(6)
0.10	0.248	0.202	0.251	0.251	1.3%
0.21	**0.241**	**0.199**	**0.251**	**0.255**	**4.0%**
0.30	0.232	0.193	0.251	0.255	7.4%
0.40	0.221	0.182	0.251	0.263	12.2%
0.50	0.207	0.169	0.251	0.277	17.4%
0.60	0.194	0.155	0.251	0.303	22.8%
0.70	0.181	0.141	0.251	0.354	28.1%
0.80	0.168	0.128	0.251	0.467	33.1%
0.90	0.156	0.116	0.251	0.822	37.8%

Note: Bold values refer to a specification with γ unconstrained and estimated as part of the optimization. The age range for both children and parents is between 30 and 40 years. Estimation are based on 336 unique parent-child pairs for children born in the 1960s and 1970s.

Table 4.9: Variances by Child-Cohort (Age: 30-40)

Variable	**Generation**	**All Cohorts** (1)	**1960s Cohort** (2)	**1970s Cohort** (3)
Earnings	Parent	0.199	0.172	0.225
	Child	0.251	0.243	0.259
Other Income	Parent	0.845	0.945	0.752
	Child	0.669	0.568	0.770
Consumption	Parent	0.097	0.112	0.081
	Child	0.118	0.100	0.135

Table 4.10: Intergenerational Elasticity Estimates by Child Cohort (Age: 30-40)

	Parameters	**All Cohorts** (1)	**1960s Cohort** (2)	**1970s Cohort** (3)
Earnings	γ	0.209	0.251	0.191
		(0.069)	(0.087)	(0.106)
Other Income	ρ	0.041	-0.006	0.099
		(0.058)	(0.068)	(0.093)
$\bar{e}_f^p$ on $n_{f,t}^k$	λ	0.217	0.202	0.244
		(0.079)	(0.131)	(0.12)
$\bar{n}_f^p$ on $e_{f,t}^k$	θ	0.040	0.009	0.079
		(0.032)	(0.046)	(0.038)
Consumption Shifters	ϕ	0.075	-0.029	0.200
		(0.075)	(0.09)	(0.124)
No. of Parent-Child Pairs	N	336	166	170

Note: Bootstrap standard errors (100 repetitions) in parentheses. Average parental ages in the three child-cohorts are 35, 36 and 35 years. Average ages of the children are 35, 34 and 35 years respectively. '*All Cohorts*' refer to the combined sample of 1960s and 1970s child birth cohorts. Food expenditure is used as proxy measure of consumption. All columns use cross-sectional data variation, net of cohort and year effects. These estimates should be compared with those in column (1) of Appendix Table C.4 where all cohorts of children are combined.

Table 4.11: Robustness: Intergenerational Elasticity Estimates

Parameters	**Baseline** (1)	$\boldsymbol{\lambda = \theta = 0}$ (2)	**Random Match** (3)	**Imputation** (4)	**Panel Data** (5)
Earnings: γ	0.230	0.340	-0.018	0.257	0.294
	(0.029)	(0.02)	(0.028)	(0.029)	(0.041)
Other Income: ρ	0.100	0.121	-0.039	0.096	0.095
	(0.029)	(0.029)	(0.025)	(0.028)	(0.045)
$\bar{e}_f^p$ on $n_{f,t}^k$: λ	0.206	0	-0.007	0.236	0.107
	(0.038)	(0)	(0.035)	(0.033)	(0.060)
$\bar{n}_f^p$ on $e_{f,t}^k$: θ	0.055	0	-0.015	0.052	0.066
	(0.017)	(0)	(0.023)	(0.015)	(0.035)
Consumption Shifters: ϕ	0.154	0.109	-0.048	0.127	0.153
	(0.034)	(0.032)	(0.034)	(0.033)	(0.046)
No. of Parent-Child Pairs: N	760	760	760	760	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. Year and cohort effects have been removed.

Table 4.12: Robustness: Idiosyncratic Components

Parameters	**Baseline** (1)	$\boldsymbol{\lambda = \theta = 0}$ (2)	**Random Match** (3)	**Imputation** (4)	**Panel Data** (5)
Parental Outcomes: Variances					
Permanent Earnings: $\sigma^2_{\bar{e}^p}$	0.295	0.289	0.291	0.291	0.289
	(0.018)	(0.025)	(0.022)	(0.02)	(0.015)
Permanent Other Income: $\sigma^2_{\bar{n}^p}$	0.806	0.806	0.808	0.807	0.478
	(0.076)	(0.074)	(0.071)	(0.072)	(0.037)
Permanent Consumption Shifters: $\sigma^2_{\bar{q}^p}$	1.031	1.053	1.032	0.861	0.689
	(0.081)	(0.08)	(0.073)	(0.07)	(0.044)
Child Idiosyncratic Heterogeneity: Variances					
Permanent Earnings: $\sigma^2_{\delta^k}$	0.228	0.214	0.247	0.224	0.208
	(0.013)	(0.014)	(0.015)	(0.011)	(0.013)
Permanent Other Income $\sigma^2_{\varepsilon^k}$	0.511	0.522	0.533	0.507	0.415
	(0.038)	(0.046)	(0.048)	(0.036)	(0.026)
Permanent Consumption Shifters: $\sigma^2_{\psi^k}$	0.730	0.741	0.752	0.573	0.584
	(0.05)	(0.063)	(0.069)	(0.04)	(0.037)
Parental Outcomes: Covariances					
Consumption Shifters & Earnings: $\sigma_{\bar{q}^p,\bar{e}^p}$	-0.271	-0.279	-0.263	-0.223	-0.258
	(0.021)	(0.029)	(0.028)	(0.019)	(0.022)
Consumption Shifters & Other Income: $\sigma_{\bar{q}^p,\bar{n}^p}$	-0.818	-0.833	-0.821	-0.769	-0.480
	(0.077)	(0.076)	(0.069)	(0.07)	(0.037)
Earnings and Other Income: $\sigma_{\bar{e}^p,\bar{n}^p}$	0.070	0.086	0.067	0.068	0.058
	(0.015)	(0.018)	(0.017)	(0.013)	(0.012)
Child Idiosyncratic Heterogeneity: Covariances					
Consumption Shifters & Earnings: σ_{ψ^k,δ^k}	-0.247	-0.253	-0.263	-0.214	-0.212
	(0.02)	(0.024)	(0.025)	(0.016)	(0.018)
Consumption Shifters & Other Income: $\sigma_{\psi^k,\varepsilon^k}$	-0.522	-0.532	-0.542	-0.480	-0.398
	(0.041)	(0.05)	(0.055)	(0.036)	(0.029)
Earnings & Other Income: $\sigma_{\delta^k,\varepsilon^k}$	0.075	0.092	0.095	0.072	0.052
	(0.013)	(0.017)	(0.019)	(0.012)	(0.013)
No. of Parent-Child Pairs: N	760	760	760	760	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. This table uses the same sample and model specification as Table 4.11.

Table 4.13: Robustness: Importance of Parents for Child Inequality

Variables	**Baseline** (1)	$\boldsymbol{\lambda = \theta = 0}$ (2)	**Random Match** (3)	**Imputation** (4)	**Panel Data** (5)
Earnings	8.0	13.5	0.1	9.4	12.3
Other Income	4.4	2.2	0.2	5.0	2.1
Consumption	29.9	19.5	0.2	47.4	22.8

Note: All numbers are percentages (%) and are based on parameter estimates in Tables 4.11 and 4.12.

4.9 Figures

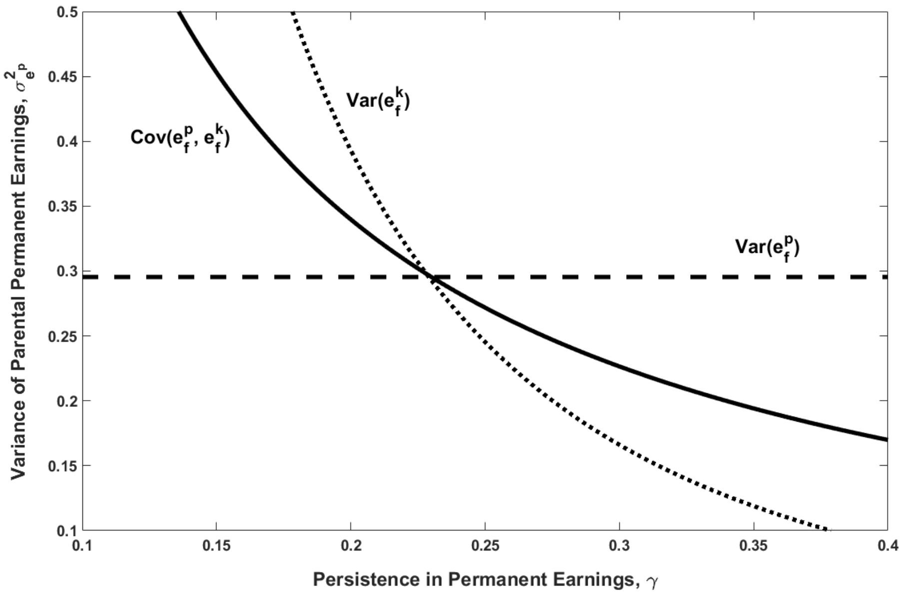

Figure 4.1: Identification of Persistence and Dispersion Parameters

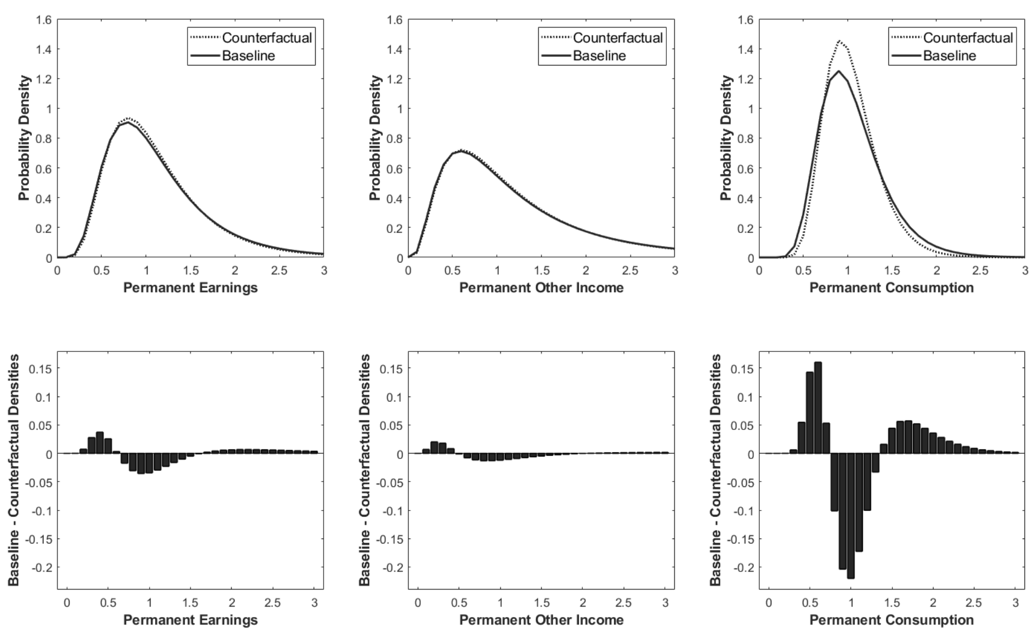

Figure 4.2: Baseline versus Counterfactual Probability Density Functions

Note: Top panels report density functions. Bottom panels report histograms of changes in local probability mass (the probability mass of the actual distribution minus the corresponding mass of the counterfactual).

Chapter 5

Conclusion

This thesis looked at the dynamics of two different types of economic choices made by firms and households. First, it studied the short-run business cycle dynamics of how the frequency of hiring and firing of workers by firms influence whether productivity increases during economic booms or busts. Second, it discussed the long-run dynamics of how economic inequality evolves across generations depending on household choices regarding educating children, intra-family transfers, consumption expenditure and so on.

Chapter 2 argued that the sudden vanishing of the procyclicality of productivity in the U.S. around the mid-1980s (the so-called productivity puzzle) is driven by a sudden reduction in the cost of hiring and firing workers. To arrive at this conclusion, it was shown that the highly procyclical factor utilization rate component of measured productivity became less important because of firms' lower reliance on adjusting the intensive margin of factors of production. Since firms typically depended on changing the intensity of factor use because of significant costs in the extensive margin of factor adjustment, lower reliance on the former indicated a drop in the cost of the latter. Moreover, it was shown that given the same positive demand shock to the U.S. economy before and after the 1980s, productivity rose much less in the post-1980 period. This indicated that firms are no longer increasing the intensity or productivity of its available workers but instead just hiring more workers to meet the extra demand, thereby corroborating the story of falling hiring and firing cost of workers. Once the drop in employment adjustment cost was established as an explanation for the productivity puzzle, I identified rapid de-unionization as a significant determinant of hiring and firing cost. In a cross-section of U.S. states and industries, it was shown that regions and sectors which experienced a larger drop in union power also underwent a bigger move towards more countercyclical productivity, and larger volatility of employment changes relative to output volatility.

While Chapter 2 showed that increased labour market flexibility due to rapid de-unionization can explain the productivity puzzle, the quantitative importance of that channel remained unexplored. This quantitative exercise was particularly revealing because simultaneous to the episode of de-unionization and the productivity puzzle in the 1980s, there was a host of other significant structural changes in the U.S. economy, and it is important to ascertain the relative explanatory power of each of these for the productivity puzzle. Chapter 3 fills precisely this gap in our understanding of the productivity puzzle. It showed through the lens of a New Keynesian model that the change in monetary policy stance by the Federal Reserve, the rising importance of technology shock vis-à-vis demand shock, the rising wage bargaining power of firms due to a decline in the power of labour unions, or the reduced volatility of shocks during Great Moderation has little to no influence in explaining the productivity puzzle. On the other hand, allowing the cost of hiring workers to fall by 67% between the pre- and post-1984 periods to reflect an equivalent drop in private-sector union membership in the U.S., was shown to be able to capture almost the entire drop in procyclicality of productivity and more than half of the rise in the relative volatility of employment fluctuations. Moving beyond the model, using a variety of industry-level datasets, it was shown that the rise of the service sector in both value-added and intermediate input use, the increased use of intangible capital, or the rising importance of sector-specific shocks vis-à-vis aggregate economy-wide shocks have no robust explanatory power for the productivity puzzle.

Chapter 4 studied how cross-sectional inequality in economic outcomes amongst any given generation of individuals get transferred to the next generation through intra-family linkages in labour earnings, other income and consumption expenditure. We first argued that inequality in the offspring's generation is determined not only by the strength of intergenerational persistence but also by the amount of inequality that prevailed in the parents' generation to begin with. We found that most of the observed inequality emanates from idiosyncratic life-cycle risk independent of parental effects. This finding is crucially determined by a moderately low (but significant) value of intergenerational persistence, which further implies that long-run inequality will not grow much through the family channel with currently observed values of persistence. Moreover, the need to jointly estimate the income and consumption processes allowing for potentially non-zero covariances between the two was also highlighted. In particular, the estimated negative covariance between permanent income and consumption shifters underlined a crucial aspect of household inequality — higher-income families consume proportionately less out of their income.

Bibliography

Aaronson, D., and B. Mazumder (2008): "Intergenerational Economic Mobility in the United States, 1940 to 2000," *Journal of Human Resources*, 43(1), 139–172.

Aaronson, D., E. Rissman, and D. G. Sullivan (2004): "Can sectoral reallocation explain the jobless recovery?," *Economic Perspectives*, (Q II), 36–39.

Aaronson, D., and D. G. Sullivan (2001): "Growth in worker quality," *Economic Perspectives*, 25(4).

Abbott, B., and G. Gallipoli (2019): "Permanent-Income Inequality," *Working Paper*.

Abbott, B., G. Gallipoli, C. Meghir, and G. L. Violante (2019): "Education Policy and Intergenerational Transfers in Equilibrium," Discussion paper, mimeo UBC, forthcoming, Journal of Political Economy.

Acemoglu, D., P. Aghion, and G. L. Violante (2001): "Deunionization, technical change and inequality," *Carnegie-Rochester Conference Series on Public Policy*, 55, 229–264.

Açıkgöz, Ö. T., and B. Kaymak (2014): "The rising skill premium and deunionization," *Journal of Monetary Economics*, 63, 37–50.

Aguiar, M., and M. Bils (2015): "Has Consumption Inequality Mirrored Income Inequality?," *American Economic Review*, 105(9), 2725–2756.

Alan, S., M. Browning, and M. Ejrnæs (2018): "Income and consumption: a micro semistructural analysis with pervasive heterogeneity," *Journal of Political Economy*, 126(5), 1827–1864.

Altonji, B. J. G., F. Hayashi, and L. J. Kotlikoff (1992): "Is the Extended Family Altruistically Linked ? Direct Tests Using Micro Data," *The American Economic Review*, 82(5), 1177–1198.

Altonji, J. G., F. Hayashi, and L. J. Kotlikoff (1997): "Parental Altruism and Inter Vivos Transfers: Theory and Evidence," *Journal of Political Economy*, 105(6), 1121–1166.

Altonji, J. G., and L. M. Segal (1996): "Small-Sample Bias in GMM Estimation of Covariance Structures," *Journal of Business & Economic Statistics*, 14(3), 353–366.

Alvarez, F., and R. Shimer (2014): "Unions and Unemployment," *Working Paper*.

ANDRESKI, P., G. LI, M. Z. SAMANCIOGLU, AND R. SCHOENI (2014): "Estimates of annual consumption expenditures and its major components in the PSID in comparison to the CE," *American Economic Review*, 104(5), 132–135.

ATTANASIO, O., AND L. PISTAFERRI (2014): "Consumption inequality over the last half century: some evidence using the new PSID consumption measure," *The American Economic Review: Papers and Proceedings*, 104(5), 122–126.

ATTANASIO, O. P., AND L. PISTAFERRI (2016): "Consumption Inequality," *Journal of Economic Perspectives*, 30(2), 3–28.

AUTOR, D. (2003): "Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing," *Journal of Labor Economics*, 21(1), 1–42.

BACHMANN, R. (2012): "Understanding the Jobless Recoveries After 1991 and 2001," *Working Paper*.

BARNICHON, R. (2010): "Productivity and unemployment over the business cycle," *Journal of Monetary Economics*, 57(8), 1013–1025.

BASU, S. (1996): "Procyclical Productivity: Increasing Returns or Cyclical Utilization?," *Quarterly Journal of Economics*, (3), 719–751.

BASU, S., J. G. FERNALD, AND M. S. KIMBALL (2001): "Are Technology Improvements Contractionary?," *The American Economic Review*, 96(1957), 1418–1448.

BAXTER, M., AND R. G. KING (1999): "Measuring Business Cycles: Approximate Band-Pass Filters for Economic Time Series," *Review of Economics and Statistics*, 81(4), 575–593.

BELLEY, P., AND L. LOCHNER (2007): "The changing role of family income and ability in determining educational achievement," *Journal of Human capital*, 1(1), 37–89.

BELLO, S. L., AND I. MORCHIO (2016): "Like Father, Like Son: Occupational Choice, Intergenerational Persistence and Misallocation," *Working Paper*.

BERGER, D. (2016): "Countercyclical Restructuring and Jobless Recoveries," *Working Paper*.

BIDDLE, J. E. (2014): "The Cyclical Behavior of Labor Productivity and the Emergence of the Labor Hoarding Concept," *Journal of Economic Perspectives*, 28(2), 197–212.

BLACK, S. E., AND P. J. DEVEREUX (2011): "Recent developments in intergenerational mobility," in *Handbook of Labor Economics*, ed. by D. Card, and O. Ashenfelter, vol. 4, pp. 773–1823. Elsevier, Amsterdam.

BLOOM, N., E. BRYNJOLFSSON, L. FOSTER, R. JARMIN, M. PATNAIK, I. SAPORTA-EKSTEN, AND J. VAN REENEN (2019): "What Drives Differences in Management Practices?," *American Economic Review*, 109(5), 1648–1683.

Blundell, R., L. Pistaferri, and I. Preston (2008): "Consumption Inequality and Partial Insurance," *The American Economic Review*, 98(5), 1887–1921.

Boar, C. (2017): "Dynastic Precautionary Savings," Discussion paper, NYU, mimeo.

Bolt, U., E. French, J. H. Maccuish, and C. ODea (2018): "Intergenerational Altruism and Transfers of Time and Money: A Life-cycle Perspective," Discussion paper, mimeo, UCL.

Bound, J., C. Brown, G. J. Duncan, and W. L. Rodgers (1994): "Evidence on the Validity of Cross-sectional and Longitudinal Labor Market Data," *Journal of Labor Economics*, 12(3), 345–368.

Brault, J., and H. Khan (2020): "The Shifts in Lead-Lag Properties of the U.S. Business Cycle," *Economic Inquiry*, 58(1), 319–334.

Bruze, G. (2016): "Intergenerational Mobility : New Evidence from Consumption Data," *Working Paper*.

Burnside, C., M. Eichenbaum, and S. Rebelo (1993): "Labor Hoarding and the Business Cycle," *Journal of Political Economy*, 101(2), 245–273.

Caballe, J. (2016): "Intergenerational mobility: measurement and the role of borrowing constraints and inherited tastes," *SERIEs*, 7(4), 393–420.

Calvo, G. A. (1983): "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, 12, 383–398.

Card, D. (1992): "Using Regional Variation in Wages to Measure the Effects of the Federal Minimum Wage," *Industrial and Labor Relations Review*, 46(1), 22–37.

Carey, M. L., and K. L. Hazelbaker (1986): "Employment growth in the temporary help industry," *Monthly Labor Review*, 109(4), 37–44.

Carneiro, P. M., and J. J. Heckman (2003): "Human Capital Policy," Discussion Paper w9495, NBER Working Paper.

Carneiro, P. M., I. Lopez Garcia, K. G. Salvanes, and E. Tominey (2015): "Intergenerational mobility and the timing of parental income," Discussion Paper 23, NHH Dept. of Economics Discussion Paper.

Cassing, S. (1996): "Correctly Measuring Real Value Added," *Review of Income and Wealth*, 42(2), 195–206.

Caucutt, E., and L. J. Lochner (2019): "Early and Late Human Capital Investments, Borrowing Constraints, and the Family," Discussion paper, forthcoming, Journal of Political Economy.

Champagne, J., and A. Kurmann (2013): "The great increase in relative wage volatility in the United States," *Journal of Monetary Economics*, 60(2), 166–183.

Champagne, J., A. Kurmann, and J. Stewart (2017): "Reconciling the divergence in aggregate U.S. wage series," *Labour Economics*, 49, 27–41.

Chang, Y., and J. H. Hong (2006): "Do Technological Improvements in the Manufacturing Sector Raise or Lower Employment," *The American Economic Review*, 96(1), 352–368.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2008): "Are structural VARs with long-run restrictions useful in developing business cycle theory?," *Journal of Monetary Economics*, 55(8), 1337–1352.

Charles, K. K., S. Danzinger, G. Li, and R. Schoeni (2014): "The Intergenerational Correlation of Consumption Expnditures," *American Economic Review*, 104(5), 136–140.

Charles, K. K., and E. Hurst (2003): "The Correlation of Wealth across Generations," *Journal of Political Economy*, 111(6), 1155–1182.

Chetty, R., N. Hendren, P. Kline, E. Saez, and N. Turner (2014): "Is the United States Still a Land of Opportunity? Recent Trends in Intergenerational Mobility," *American Economic Review: Papers & Proceedings*, 104(5), 141–147.

Christiano, L. J., M. Eichenbaum, and R. Vigfusson (2003): "What Happens After a Technology Shock?," *NBER Working Paper*, July(9819).

Christiano, L. J., and T. J. Fitzgerald (2003): "The Band Pass Filter," *International Economic Review*, 44(2), 435–465.

Comerford, D., J. V. Rodriguez Mora, and M. J. Watts (2017): "The rise of meritocracy and the inheritance of advantage," Discussion paper, University of Strathclyde.

Corak, M. (2013): "Income Inequality, Equality of Opportunity, and Intergenerational Mobility," Discussion Paper No. 7520, IZA Discussion Paper Series.

Corak, M., and P. Piraino (2010): "Intergenerational Earnings Mobility and the Inheritance of Employers," Discussion Paper 4876, IZA Discussion Paper Series.

Cordoba, J. C., X. Liu, and M. Ripoll (2016): "Fertility, social mobility and long run inequality," *Journal of Monetary Economics*, 77, 103–124.

Daruich, D., S. D. Addario, and R. Saggio (2017): "The Effects of Partial Employment Protection Reforms: Evidence from Italy," *Working Paper*.

Daruich, D., and J. Kozlowski (2016): "Explaining Income Inequality and Intergenerational Mobility: The Role of Fertility and Family Transfers," Discussion paper, Society for Economic Dynamics Meetings.

Davis, S. J., R. J. Faberman, and J. Haltiwanger (2006): "The Flow Approach to Labor Markets: New Data Sources and Micro-Macro Links," *Journal of Economic Perspectives*, 20(3), 3–26.

Denton, F. T. (1971): "Adjustment of Monthly or Quarterly Series to Annual Totals: An Approach Based on Quadratic Minimization," *Journal of the American Statistical Association*, 66(333), 99–102.

Dinlersoz, E., and J. Greenwood (2016): "The rise and fall of unions in the United States," *Journal of Monetary Economics*, 83, 129–146.

Doppelt, R., and K. O'Hara (2018): "Bayesian Estimation of Fractionally Integrated Vector Autoregressions and an Application to Identified Technology Shocks," *Working Paper*.

Dornbusch, R., and S. Fischer (1981): *Macroeconomics*. New York: McGraw Hill.

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004): "Do the Rich Save More?," *Journal of Political Economy*, 112(2), 397–444.

Elsby, M. W., B. Hobijn, and A. Sahin (2015): "On the importance of the participation margin for labor market fluctuations," *Journal of Monetary Economics*, 72, 64–82.

Fan, S. C. (2006): "Do the Rich Save More? A New View Based on Intergenerational Transfers," *Southern Economic Journal*, 73(2), 362–373.

Farber, H. S., and B. Western (2002): "Ronald Reagan and the Politics of Declining Union Organization," *British Journal of Industrial Relations*, 40(3), 385–401.

Felix, R. A., and J. R. Hines Jr. (2009): "Corporate Taxes and Union Wages in the United States," *NBER Working Paper Series*, 15263(August).

Fernald, J. (2014): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," *Federal Reserve Bank of San Francisco Working Paper Series*, April(19).

Fernandez, R., A. Fogli, and C. Olivetti (2004): "Mothers and Sons : Preference Formation and Female Labor Force Dynamics," *The Quarterly Journal of Economics*, 119(4), 1249–1299.

Ferraresi, T., A. Roventini, and W. Semmler (2016): "Macroeconomic Regimes, Technological Shocks and Employment Dynamics," *Working Paper*.

Fisher, H., and H. Low (2015): "Financial implications of relationship breakdown: Does marriage matter?," *Review of Economics of the Household*, 13(4), 735–769.

Flavin, M., and T. Yamashita (2002): "Owner-Occupied Housing and the Composition of the Household Portfolio," *The American Economic Review*, 92(1), 345–362.

Foerster, A. T., P.-D. G. Sarte, and M. W. Watson (2011): "Sectoral versus Aggregate Shocks: A Structural Factor Analysis of Industrial Production," *Journal of Political Economy*, 119(1), 1–38.

Foll, T., and A. Hartmann (2019): "A Joint Theory of Polarization and Deunionization," *Working Paper*.

Foster, L., C. Grim, and J. Haltiwanger (2016): "Reallocation in the Great Recession: Cleansing or Not?," *Journal of Labor Economics*, 34(S1), S293–S331.

Francis, N., M. T. Owyang, J. E. Rousch, and R. DiCecio (2014): "Technology Diffusion and Productivity Growth in Health Care," *Review of Economics and Statistics*, 96(4), 638–647.

Freeman, R. B., and J. L. Medoff (1982): "Substitution Between Production Labor and Other Inputs in Unionized and Nonunionized Manufacturing," *The Review of Economics and Statistics*, 64(2), 220–233.

Galí, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *The American Economic Review*, 89(1), 249–271.

——— (2008): *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*. Princeton University Press.

——— (2011): "Monetary Policy and Unemployment," *Handbook of Monetary Economics*, 3A, 487–546.

Galí, J., and L. Gambetti (2009): "On the Sources of the Great Moderation," *American Economic Journal: Macroeconomics*, 1(1), 26–57.

Galí, J., and T. van Rens (2017): "The Vanishing Procyclicality of Labor Productivity," *Working Paper*.

Garin, J., M. J. Pries, and E. R. Sims (2018): "The Relative Importance of Aggregate and Sectoral Shocks and the Changing Nature of Economic Fluctuations," *American Economic Journal: Macroeconomics*, 10(1), 119–148.

Gayle, G.-L., L. Golan, and M. A. Soytas (2018): "What is the Source of the Intergenerational Correlation in Earnings?," *Working Paper*.

Giupponi, G., and C. Landais (2018): "Subsidizing Labor Hoarding in Recessions: The Employment and Welfare Effects of Short Time Work," *London School of Economics Working Paper*.

Gnocchi, S., and E. Pappa (2009): "Do labor market rigidities matter for business cycles? Yes they do," *Barcelona Economics Working Paper Series*, July(411).

Gordon, R. J. (1993): "The Jobless Recovery: Does It Signal a New Era of Productivity-led Growth?," *Brookings Papers on Economic Activity*, 1, 271–316.

——— (2011): "The Evolution of Okun's Law and of Cyclical Productivity Fluctuations in the United States and in the EU-15," in *EES/IAB Workshop, Labor Market Institutions and the Macroeconomy.*

Gouskova, E., N. Chiteji, and F. Stafford (2010): "Estimating the intergenerational persistence of lifetime earnings with life course matching: Evidence from the PSID," *Labour Economics*, 17(3), 592–597.

Graves, S. (2019): "The State Dependent Effectiveness of Hiring Subsidies," *Job Market Paper.*

Grawe, N. D. (2006): "Lifecycle bias in estimates of intergenerational earnings persistence," *Labour Economics*, 13(5), 551–570.

Groshen, E. L., and S. Potter (2003): "Has Structural Change Contributed to a Jobless Recovery?," *Current Issues in Economics & Finance*, 9(8), 1–7.

Gu, G. W., and E. Prasad (2018): "New Evidence on Cyclical Variation in Labor Costs in the U.S.," *NBER Working Paper.*

Hagedorn, M., and I. Manovskii (2008): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited," *American Economic Review*, 98(4), 1692–1706.

Haider, S., and G. Solon (2006): "Life-cycle variation in the association between current and lifetime earnings," *American Economic Review*, 96(4), 1308–1320.

Hamermesh, D., and A. Rees (1984): *The Economics of Work and Pay.* New York: Harper and Rowe.

Heathcote, J., F. Perri, and G. L. Violante (2010): "Unequal we stand: An empirical analysis of economic inequality in the United States, 1967 – 2006," *Review of Economic Dynamics*, 13(1), 15–51.

Heckman, J., C. Pagés-Serra, A. C. Edwards, and P. Guidotti (2000): "The Cost of Job Security Regulation: Evidence from Latin American Labor Markets," *Economía*, 1(1), 109–154.

Herrendorf, B., C. Herrington, and Á. Valentinyi (2015): "Sectoral Technology and Structural Transformation," *American Economic Journal: Macroeconomics*, 7(4), 104–133.

Hertz, T. (2007): "Trends in the intergenerational elasticity of family income in the United States," *Industrial Relations*, 46(1), 22–50.

Hirsch, B. T., and D. A. Macpherson (2003): "Union Membership and Coverage Files from the Current Population Surveys: Note," *Industrial and Labor Relations Review*, 56(2), 349–354.

Hodrick, R. J., and E. C. Prescott (1997): "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking*, 29(1), 1–16.

Holmlund, H., M. Lindahl, and E. Plug (2011): "The causal effect of parents' schooling on children's schooling: A comparison of estimation methods," *Journal of Economic Literature*, 49(3), 615–51.

Hryshko, D., and I. Manovskii (2019): "How much consumption insurance in the U.S.?," Discussion paper, Univ. of Alberta and Univ. of Pennsylvania.

Hsieh, C.-T., and P. J. Klenow (2009): "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, CXXIV(4), 1403–1448.

Huang, K. X., Z. Liu, and L. Phaneuf (2004): "Why does the cyclical behavior of real wages change over time?," *American Economic Review*, 94(4), 836–856.

Iacobucci, A., and A. Noullez (2005): "A Frequency Selective Filter for Short-Length Time Series," *Computational Economics*, 25(1-2), 75–102.

Jaimovich, N., and H. E. Siu (2015): "Job Polarization and Jobless Recoveries," *Working Paper*.

Jalón, B., S. Sosvilla-Rivero, and J. A. Herce (2017): "Countercyclical Labor Productivity: The Spanish Anomaly," *Research Institute of Applied Economics Working Paper*, 12, 1–29.

Jorda, Ò. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *The American Economic Review*, 95(1), 161–182.

Jorgenson, D. W., M. S. Ho, and J. D. Samuels (2012): "A Prototype Industry-Level Production Account for the United States, 1947-2010," *Second World KLEMS Conference*, August.

Kanbur, R. (2019): "In Praise of Snapshots," *CEPR Discussion Paper Series*, DP(14093).

Kaplan, G., and G. L. Violante (2010): "How much consumption insurance beyond self-insurance?," *American Economic Journal: Macroeconomics*, 2(4), 53–87.

Krueger, A. B. (2012): "The Rise and Consequences of Inequality in the United States," *Discussion paper, Center for American Progress, Presentation made on January 12th.*

Krueger, D., and F. Perri (2006): "Does Income Inequality Lead to Consumption Inequality? Evidence and Theory," *The Review of Economic Studies*, 73(1), 163–193.

Krusell, P., L. E. Ohanian, J.-V. Rios-Rull, and G. L. Violante (2000): "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68(5), 1029–1053.

Landersø, R., and J. J. Heckman (2017): "The Scandinavian Fantasy: The Sources of Intergenerational Mobility in Denmark and the US," *Scandinavian Journal of Economics*, 119(1), 178–230.

Lazear, E. P., K. L. Shaw, and C. Stanton (2016): "Making Do with Less: Working Harder During Recessions," *Journal of Labor Economics*, 34(1), S333–S360.

Lee, C.-I., and G. Solon (2009): "Trends in Intergenerational Income Mobility," *Review of Economics and Statistics*, 91(4), 766–772.

Lee, S. Y., N. Roys, and A. Seshadri (2014): "The causal effect of parental human capital on childrens human capital," *Unpublished manuscript, University of Wisconsin–Madison, Department of Economics.*

Lee, S. Y. T., and A. Seshadri (2019): "On the Intergenerational Transmission of Economic Status," *Journal of Political Economy*, 127(2), 855–921.

Lefgren, L., D. Sims, and M. J. Lindquist (2012): "Rich dad, smart dad: Decomposing the intergenerational transmission of income," *Journal of Political Economy*, 120(2), 268–303.

Lemieux, T., W. B. Macleod, and D. Parent (2009): "Performance Pay and Wage Inequality," *The Quarterly Journal of Economics*, 124(1), 1–49.

Li, G., R. F. Schoeni, S. Danziger, and K. K. Charles (2010): "New expenditure data in the PSID: comparisons with the CE," *Monthly Labor Review*, February, 29–39.

Mayer, S. E., and L. M. Lopoo (2005): "On the Intergenerational Transmission of Economic Status," *The Journal of Human Resources*, XL(1), 169–185.

McGrattan, E. R., and E. C. Prescott (2012): "The Labor Productivity Puzzle," *Federal Reserve Bank of Minneapolis Research Department Working Paper*, May(694).

Merz, M., and E. Yashiv (2007): "Labor and the Market Value of the Firm," *The American Economic Review*, 97(4), 1419–1431.

Molnarova, Z. (2020): "Industry evidence and the vanishing cyclicality of labor productivity," *University of Vienna Working Papers.*

Mortensen, D. T., and É. Nagypál (2007): "More on unemployment and vacancy fluctuations," *Review of Economic Dynamics*, 10, 327–347.

Mulligan, C. B. (2011): "Rising Labor Productivity during the 2008-9 Recession," *NBER Working Paper*, November(02138).

Murali, S. (2018): "Labor Productivity Puzzle: An Explanation Using Empirical Mode Decomposition," in *Essays in Macroeconomics*, pp. 72–88.

Nucci, F., and M. Riggi (2013): "Performance pay and changes in U.S. labor market dynamics," *Journal of Economic Dynamics and Control*, 37(12), 2796–2813.

Ohanian, L. E., and A. Raffo (2012): "Aggregate hours worked in OECD countries: New measurement and implications for business cycles," *Journal of Monetary Economics*, 59(1), 40–56.

Okun, A. M. (1962): "Potential GNP: Its Measurement and Significance," *Proceedings of the Business and Economic Statistics Section of the American Statistical Association.*

Oreopoulos, P., M. E. Page, and A. H. Stevens (2006): "The intergenerational effects of compulsory schooling," *Journal of Labor Economics*, 24(4), 729–760.

Panovska, I. B. (2017): "What Explains the Recent Jobless Recoveries?," *Macroeconomic Dynamics*, 21(3), 708–732.

Peters, H. E. (1992): "Patterns of Intergenerational Mobility in Income and Earnings," *The Review of Economics and Statistics*, 74(3), 456–466.

Ramey, V. A. (2016): "Macroeconomic Shocks and Their Propagation," *Handbook of Macroeconomics*, pp. 71–162.

Rauh, C. (2017): "Voting, education, and the Great Gatsby Curve," *Journal of Public Economics*, 146, 1–14.

Ravn, M. O., and H. Uhlig (2002): "On adjusting the Hodrick-Prescott filter for the frequency of observations," *Review of Economics and Statistics*, 84(2), 371–380.

Restuccia, D., and C. Urrutia (2004): "Intergenerational persistence of earnings: The role of early and college education," *American Economic Review*, 94(5), 1354–1378.

Riddell, C. W. (1993): *Unionization in Canada and the United States: A Tale of Two Countries*, no. January.

Rissman, E. R. (1997): "Measuring labor market turbulence," *Economic Perspectives*, 21(3), 2–14.

Romer, C. D., and D. H. Romer (2004): "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 94(4), 1055–1084.

Schreft, L. S., and A. Singh (2003): "A Closer Look at Jobless Recoveries," *Economic Review*, (Q II), 45–73.

Shimer, R. (2012): "Reassessing the ins and outs of unemployment," *Review of Economic Dynamics*, 15(2), 127–148.

Shorrocks, A. F. (1978): "The Measurement of Mobility," *Econometrica*, 46(5), 1013–1024.

Silva, J. I., and M. Toledo (2009): "Labor Turnover Costs and the Cyclical Behavior of Vacancies and Unemployment," *Macroeconomic Dynamics*, 13(S1), 76–96.

Smets, F., and R. Wouters (2007): "Shocks and Frictions in US Business Cycles : A Bayesian DSGE Approach," *The American Economic Review*, 97(3), 586–606.

Soskice, D., and T. Iversen (2000): "The Nonneutrality of Monetary Policy," *The Quarterly Journal of Economics*, 115(1), 265–284.

Stansbury, A., and L. H. Summers (2020): "The Declining Worker Power Hypothesis: An Explanation for the Recent Evolution of the American Economy," *NBER Working Paper Series*, May(27193).

Straub, L. (2018): "Consumption, Savings, and the Distribution of Permanent Income," *Working Paper*.

Taschereau-Dumouchel, M. (2017): "The Union Threat," *Working Paper*.

Troy, L., and N. Sheflin (1985): *U.S. Union Source Book: Membership, Finances, Structures, Directory*. West Orange, NJ: IRDIS.

Van Zandweghe, W. (2010): "Why Have the Dynamics of Labor Productivity Changed?," *Economic Review*, (Q III), 5–30.

Waldkirch, A., S. Ng, and D. Cox (2004): "Intergenerational Linkages in Consumption Behavior," *Journal of Human Resources*, 39(November 2002), 355–381.

Wang, J. C. (2014): "Vanishing Procyclicality of Productivity? Industry Evidence," *Federal Reserve Bank of Boston Working Papers*, (14-15).

Zanetti, F. (2007): "A non-Walrasian labor market in a monetary model of the business cycle," *Journal of Economic Dynamics and Control*, 31, 2413–2437.

Appendix A

Appendix to Chapter 2

A.1 Robustness to Choice of Filters and Datasets

In this section, I will present the cyclical correlations and volatilities of different variables using different datasets and time-series filters. In particular, the three datasets considered here are as follows: (i) Labor Productivity and Costs (LPC) dataset published by the Bureau of Labor Statistics (BLS) that contains both quarterly and annual data on output, hours, employment and labour productivity for the U.S. business sector; (ii) John G. Fernald's TFP dataset which contains quarterly and annual data on growth rates of TFP, factor utilization rate and utilization-adjusted TFP for the U.S. business sector; and (iii) KLEMS dataset (compiled by Jorgenson, Ho and Samuels) that contains annual data on ouput, hours, employment, labour productivity and growth rate of TFP for the aggregate U.S. economy.

Table A.1: Cyclical Correlation of Average Labour Productivity (Output per Hour)

	With Output			With Hours			With Employment		
Dataset & Filter Choice	Pre 1983	Post 1984	Change	Pre 1983	Post 1984	Change	Pre 1983	Post 1984	Change
Panel A: LPC Data									
Hodrick-Prescott (λ=1600)	0.61	-0.01	-0.62	0.15	-0.53	-0.68	0.05	-0.59	-0.64
BK-Bandpass: 6-32 Qtrs.	0.56	-0.03	-0.59	0.12	-0.53	-0.65	0.01	-0.58	-0.59
Quarterly Growth Rate	0.71	0.53	-0.18	0.02	-0.34	-0.36	-0.02	-0.33	-0.31
4-Quarter Growth Rate	0.63	0.23	-0.40	0.08	-0.37	-0.45	-0.04	-0.37	-0.34
Annual Growth Rate	0.64	0.16	-0.48	0.12	-0.40	-0.52	-0.03	-0.40	-0.37
Panel B: KLEMS Data									
Hodrick-Prescott (λ=6.25)	0.35	-0.02	-0.37	-0.22	-0.62	-0.40	-0.28	-0.60	-0.32
BK-Bandpass: 2-8 Years	0.42	0.32	-0.10	-0.17	-0.52	-0.35	-0.33	-0.42	-0.10
Annual Growth Rate	0.53	0.22	-0.31	-0.10	-0.32	-0.22	-0.14	-0.24	-0.10

Since the TFP data is only available in growth rates, I could only use quarterly and annual growth rates as the filter for the analysis involving TFP. Apart from growth rates, I have considered two other time-series filters that are regularly used for extracting business cycle dynamics from macro-data — (i) Hodrick and Prescott (1997) filter, with the smoothing parameter being 1600 for quarterly data and 6.25 for annual data, following Ravn and Uhlig (2002), and (ii) bandpass filter, extracting the dynamics between 6 and 32 quarters or between 2 and 8 years for quarterly or annual data respectively. To compare the two sub-periods — before and after the mid-1980s, I have presented all the moments in Tables A.1 through A.3 separately, and also indicated the statistical significance of the difference between the two sub-periods.

Note that in these tables there are two choices for the bandpass filter — (i) the Baxter and King (1999) (BK) filter, and (ii) the Christiano and Fitzgerald (2003) (CF) filter. I use the BK filter for any analysis involving correlations. This is because the BK filter, unlike the CF filter, does not introduce any time- or frequency-dependent phase shift in the filtered data (see Iacobucci and Noullez (2005)). While using the CF filter might introduce spurious correlations in the filtered data, the BK filter distorts the amplitude or volatility of the extracted cycle. This prompts me to use the CF filter for the analysis involving cyclical volatility.

Table A.2: Cyclical Volatility of Output, Hours & Employment

	s.d.(Output)			**s.d.(Hours)**			**s.d.(Employment)**		
Dataset & Filter Choice	Pre-1983	Post-1984	$\frac{Post}{Pre}$	Pre-1983	Post-1984	$\frac{Post}{Pre}$	Pre-1983	Post-1984	$\frac{Post}{Pre}$
Panel A: LPC Data									
Hodrick-Prescott (λ=1600)	2.42	1.41	0.58	1.95	1.66	0.80	1.61	1.38	0.85
CF-Bandpass: 6-32 Quarters	2.33	1.36	0.58	1.88	1.46	0.78	1.53	1.14	0.74
4-Quarter Growth Rate	0.94	0.59	0.63	0.71	0.60	0.85	0.60	0.50	0.84
Panel B: KLEMS Data									
Hodrick-Prescott (λ=6.25)	1.68	0.94	0.56	1.59	1.18	0.74	1.38	0.91	0.66
CF-Bandpass: 2-8 Years	1.65	1.02	0.62	1.56	1.13	0.72	1.34	0.85	0.63
Annual Growth Rate	2.73	1.89	0.69	2.28	1.91	0.84	1.99	1.63	0.82

Table A.3: Relative Cyclical Volatility of Hours & Employment

	$\frac{s.d.(Hours)}{s.d.(Output)}$			$\frac{s.d.(Employment)}{s.d.(Output)}$			$\frac{s.d.(Employment)}{s.d.(Hours/Worker)}$		
Dataset & Filter Choice	Pre-1983	Post-1984	$\frac{Post}{Pre}$	Pre-1983	Post-1984	$\frac{Post}{Pre}$	Pre-1983	Post-1984	$\frac{Post}{Pre}$
Panel A: LPC Data									
Hodrick-Prescott (λ=1600)	0.80	1.18	1.47	0.67	0.98	1.46	2.99	3.17	1.06
CF-Bandpass: 6-32 Quarters	0.81	1.08	1.33	0.66	0.84	1.28	3.13	2.71	0.87
4-Quarter Growth Rate	0.76	1.02	1.35	0.64	0.85	1.34	2.82	3.14	1.11
Panel B: KLEMS Data									
Hodrick-Prescott (λ=6.25)	0.95	1.26	1.33	0.82	0.97	1.19	3.50	2.73	0.78
CF-Bandpass: 2-8 Years	0.95	1.11	1.17	0.82	0.83	1.02	3.28	2.47	0.75
Annual Growth Rate	0.83	1.01	1.22	0.73	0.86	1.19	3.24	3.47	1.07

In Table A.4, I show that the vanishing procyclicality volatility reduction of factor utilization rate is not unique to using the hours per worker proxy used in Fernald (2014). The Capacity Utilization Rate published by the Federal Reserve Board (FRB) based on the Quarterly Survey of Plant Capacity (QSPC) by the Census Bureau paints a very similar story. The QSPC asks plants to report both their current production and their full production capacity, defined as "*the maximum level of production that this establishment could reasonably expect to attain under normal and realistic operating conditions fully utilizing the machinery and equipment in place*". While hours per worker and the capacity utilization survey measure very different quantities, the correlation between the growth rates of factor utilization and capacity utilization rates is 0.73.

Table A.4: Reduction in Procyclicality and Volatility of Factor/Capacity Utilization Rates

	Corr. with Output			**Corr. with Hours**			**Variance**		
Utilization Rates	Pre 1983	Post 1984	Change	Pre 1983	Post 1984	Change	Pre 1983	Post 1984	Change
Factor (Fernald)	0.73	0.49	- 0.24	0.67	0.52	- 0.15	11.67	1.64	-85.9%
Capacity (FRB)	0.86	0.61	- 0.25	0.89	0.64	- 0.25	8.28	4.73	-42.9%

A.2 Evidence for De-unionization: A Difference-in-Difference Strategy

I will use sectoral variation across U.S. industries to see if de-unionization caused labour productivity correlation to fall. To argue for this causal channel, I follow a difference-in-difference regression strategy similar to Card (1992). I consider a very simple structural model that explains the fall in employment adjustment cost in industry i, $\Delta Cost_i$, as a function of the fraction of workers unionized in the industry prior to mid-1980s, $Union_i^{pre}$, and the change in correlation of labour productivity with hours worked, $\Delta Corr\,(lp_i, h_i)$, as a function of that change in cost:

$$\Delta Cost_i = a + bUnion_i^{pre} + e_i \tag{A.1}$$

$$\Delta Corr\,(lp_i, h_i) = \alpha + \beta \Delta Cost_i + \varepsilon_i \tag{A.2}$$

The above system of structural equations can be combined to a reduced-form correlation change equation:

$$\begin{aligned} \Delta Corr\,(lp_i, h_i) &= (\alpha + a\beta) + b\beta Union_i^{pre} + (\beta e_i + \varepsilon_i) \\ \implies \Delta Corr\,(lp_i, h_i) &\equiv \beta_0 + \beta_1 Union_i^{pre} + \eta_i \end{aligned} \tag{A.3}$$

Equation (A.3) can be interpreted as showing the impact on productivity correlations in different industries which were differentially impacted by de-unionization. In other words, if one thinks of the fall in union rates around the early 1980s as the treatment, then the intensity of treatment varied across industries according to the pre-intervention level of union densities in those industries. In particular, an industry with a higher pre-intervention level of union density should be impacted more by the de-unionization treatment, thereby leading to a larger fall in productivity correlations. As an extreme example, an industry with no unionization to begin with will experience no impact of the de-unionization event. Running the regression in equation (A.3) across 17 U.S. industries, I find a significant positive effect of union density on the fall in productivity correlation, as shown in Figure A.1. In order to avoid small industries driving the correlation pattern, I weighted the observations by the pre-1983 average industry employment level.

Finally, replacing the change in productivity correlations by the change in the relative volatility of employment in equation (A.3), I find that industries with a larger pre-1984 level of union density experienced a larger increase (or a smaller decrease) in the volatility of employment relative to that of output and hours per worker. This is shown in Figure A.2.

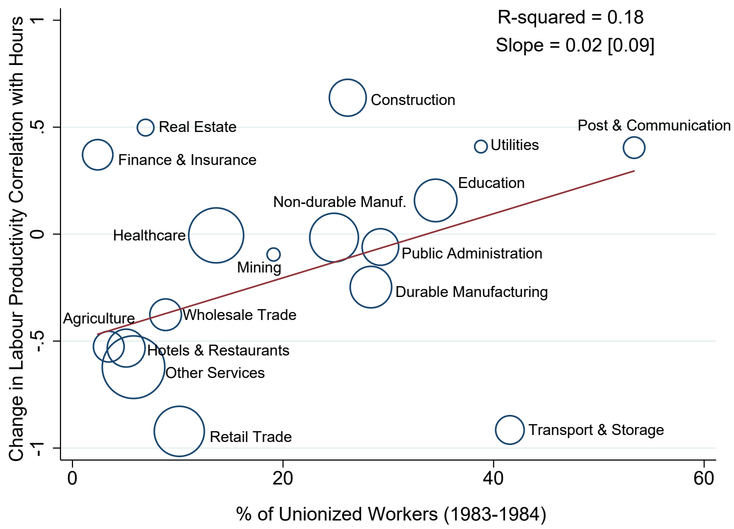

Figure A.1: Difference-in-Difference Effect of Union Density on Productivity Correlation

Note: Data on industry-level unionization rates comes from the Current Population Survey (CPS), collected by Hirsch and Macpherson (2003). Labour productivity is defined as real value added per hour worked. Data on value-added, hours and employment comes from KLEMS dataset. CPS industry codes for unionization and SIC industry codes for labour productivity were matched to create a consistent set of 17 U.S. industries. The Baxter and King (1999) bandpass filter between 2 and 8 years have been used to de-trend the variables. Change in productivity correlations is the difference in correlation between the post-1984 period (1984-2003) and the pre-1983 period (1964-1983). Since industry-level union data is available only from 1983 onwards, I have used the 2-year average of 1983 and 1984 values as the measure of pre-1984 level of union density. Size of the bubbles represent pre-1983 average industry employment level. The p-value of the slope coefficient using robust standard errors is reported in parentheses.

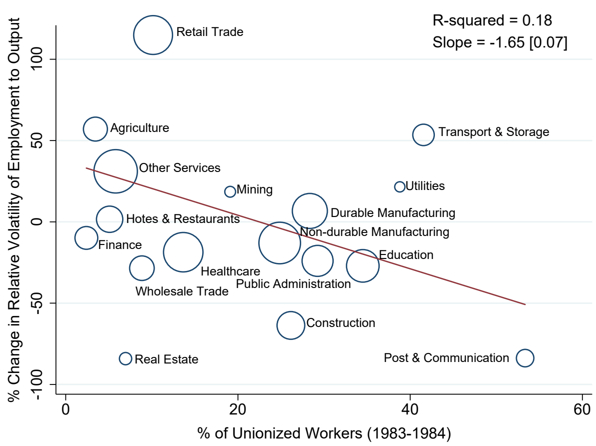

Relative to Output Volatility

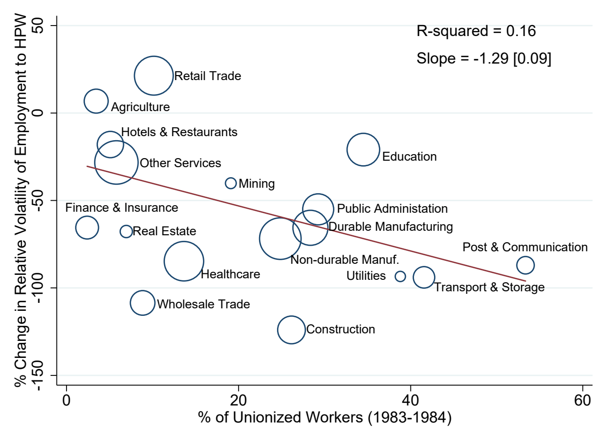

Relative to Hours Per Worker Volatility

Figure A.2: Difference-in-Difference Effect of Union Density on Relative Volatility of Employment

Note: Data on industry-level unionization rates comes from the Current Population Survey (CPS), collected by Hirsch and Macpherson (2003). Labour productivity is defined as real value added per hour worked. Data on value-added, hours and employment comes from KLEMS dataset. CPS industry codes for unionization and SIC industry codes for labour productivity were matched to create a consistent set of 17 U.S. industries. The Baxter and King (1999) bandpass filter between 2 and 8 years have been used to de-trend the variables. Change in productivity correlations is the difference in correlation between the post-1984 period (1984-2003) and the pre-1983 period (1964-1983). Since industry-level union data is available only from 1983 onwards, I have used the 2-year average of 1983 and 1984 values as the measure of pre-1984 level of union density. Size of the bubbles represent pre-1983 average industry employment level. The p-value of the slope coefficient using robust standard errors is reported in parentheses.

A.3 Choice of SVAR Specification

The seminal paper of Galí (1999) showed that labour input responds negatively to technology shocks on impact. In Galí's Vector Auto-Regression (VAR) specification, technology shocks were identified as the only shock that could change productivity in the long run.[87] Since this finding was at odds with the standard wisdom of a real business cycle model where technology shocks are positively correlated with both output and hours input, a lot of criticism was generated against this finding.

The main criticism of Galí's finding was that it was not robust to how the variables in the VAR, particularly the measure of labour input, were filtered.[88]Christiano, Eichenbaum, and Vigfusson (2003) show that filtering the measure of labour inputs by taking its growth rate generates the spurious negative impulse response of per capita hours to a positive technology shock. They argue that per capita hours worked cannot be a non-stationary process, and hence differencing an already stationary time series creates the spurious negative correlation. In fact, when per capita hours enters the SVAR in levels, instead of growth rates, technology shocks indeed become positively correlated with hours. Nevertheless, it has since been argued that not controlling for low-frequency movements in the labour input might introduce spurious correlations with productivity growth. A host of new VAR estimation techniques, like Threshold VAR by Ferraresi, Roventini, and Semmler (2016), and Bayesian estimation of Fractionally Integrated VAR by Doppelt and O'Hara (2018) — all corroborate that after controlling for low-frequency movements, hours per capita responds negatively to a technology shock on impact.

I will use the technique in Galí and Gambetti (2009) to control for the low-frequency movements in per capita hours worked, and use the same identifying assumption as in Galí (1999). Galí and Gambetti (2009) use a VAR model with time-varying coefficients and stochastic volatility of the innovations. Defining $x_t \equiv [\Delta (y_t - n_t), n_t]$, where y_t and n_t denote the (log) output and (log) hours in per capita terms, the reduced form VAR can be written as:

$$x_t = A_{0,t} + A_{1,t}x_{t-1} + A_{2,t}x_{t-2} + ... + A_{p,t}x_{t-p} + u_t \tag{A.4}$$

where $A_{0,t}$ is a vector of time-varying intercepts, $A_{i,t}$, $i = 1, ..., p$ are matrices of time-varying coefficients, and the sequence of innovations $\{u_t\}$ follows a Gaussian white noise process (uncorrelated with all lages of x_t) with zero mean and time-varying covariance matrix. Crucially, the presence of a

[87] In a two-variable SVAR with productivity growth and per capita hours, the identifying assumption implies that the long run coefficient matrix is lower triangular, that is, $\begin{pmatrix} \Delta (y_t - n_t) \\ n_t \end{pmatrix} = \begin{pmatrix} C_{11}(L) & 0 \\ C_{21}(L) & C_{22}(L) \end{pmatrix} \begin{pmatrix} \varepsilon_t^a \\ \varepsilon_t^\nu \end{pmatrix}$, where ε_t^a is the technology shock, and ε_t^ν is the non-technology or demand shock.

[88] There were other criticisms as well. For example, Chang and Hong (2006) argue that TFP should be used instead of labour productivity as the measure of productivity in the VAR to properly identify technological shocks. I have compared the results obtained by using both these measures of productivity in details. Chari, Kehoe, and McGrattan (2008) argue that the use of long run restrictions in structural VAR to identify shocks, like Galí's identification argument, is not helpful for developing business cycle theories in general. However, Francis, Owyang, Rousch, and DiCecio (2014) provide a flexible finite-horizon alternative to the long run restrictions, and corroborate Galí's conclusions.

time-varying intercept in equation A.4 absorbs the low-frequency co-movement between productivity growth and per capita hours, thereby overcoming potential distortions in the VAR estimation. There are two main advantages of this specification: first, it allows one to control for low-frequency movements in per capita hours without having to extract the cyclical component of hours through any form of ad hoc time series filtering, and second, it allows one to know the complete dynamics of the impulse responses over the years so that it can be pin-pointed as to exactly when the responses began to change. Nonetheless, this method of controlling for the low-frequency movements in per capita hours also generates a negative response of hours to a positive technology shock.

As an alterantive to VAR specifications, which require strong identifying assumptions, I present an alternative methodology, à la Jorda (2005), of estimating the impulse response of hours to changes in utilization-adjusted TFP. For this projection-type analysis, I run the regression specification used by Ramey (2016):

$$\ln\left(hours_t/pop_t\right) = \alpha_h + \beta_h \Delta \ln\left(uatfp_t\right) + \theta_h\left(L\right) X_{t-1} + \varepsilon_{t+h} \tag{A.5}$$

β_h: Response of hours at time $t+h$ to a technology shock at time t.
X_{t-1}: One-period lagged values of growth rate of utilization-adjusted TFP ($uatfp$), log per capita hours, log real GDP per capita, log labour productivity, and log real stock prices per capita.
ε_{t+h} is serially correlated, and so standard errors incorporate Newey-West correction.

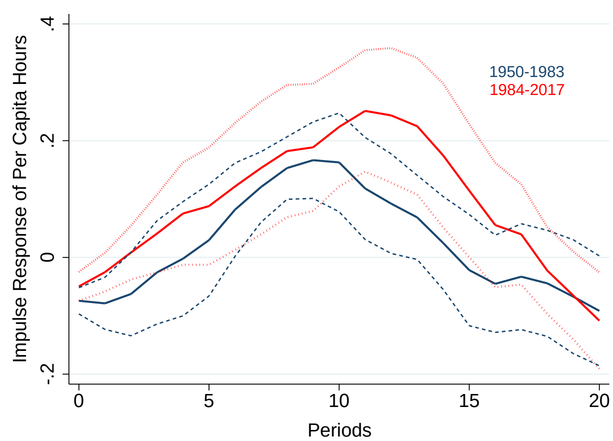

Figure A.3: IRF of Per Capita Hours to Utilization-Adjusted TFP Shock

Note: The solid blue and red lines are the impulse responses of per capita hours to one percent rise in utilization-adjusted TFP in the pre-1983 and post-1984 periods respectively. The corresponding dashed and dotted lines are the 90 percent confidence intervals for the impulse responses. All data for the regression come from Ramey (2016).

This methodology of a simple regression model with the shock being the explanatory variable not only shows the negative correlation of hours and technology shock but also that the negative response of hours became muted after the mid-1980s (see Figure A.3.)

A.4 Impulse Response Functions from Time-Varying SVARs

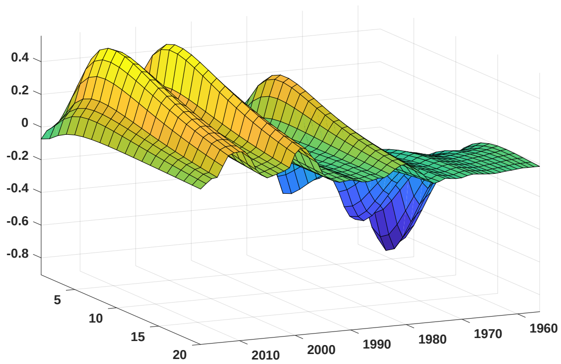

(a) Technology Shock: Hours

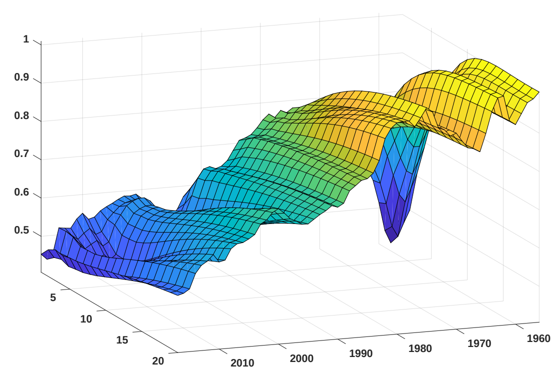

(b) Technology Shock: Labour Productivity

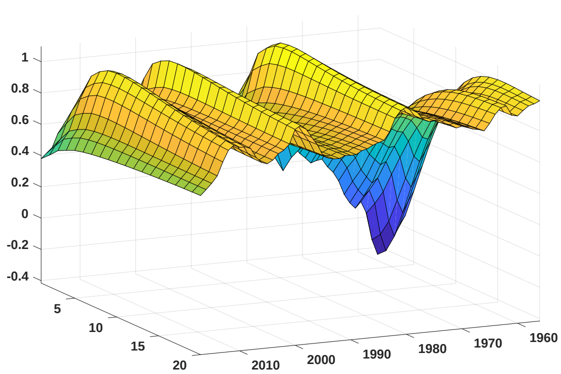

(c) Technology Shock: Output

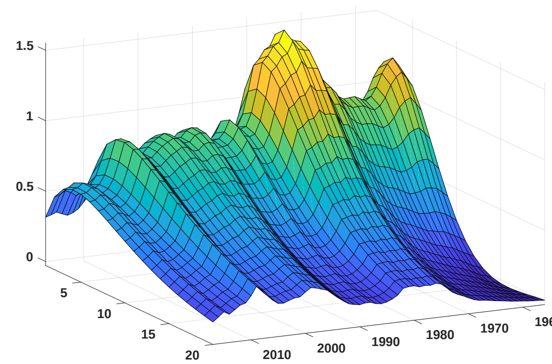

(d) Demand Shock: Hours

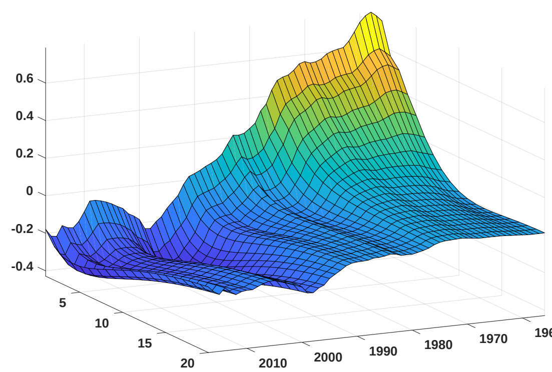

(e) Demand Shock: Labour Productivity

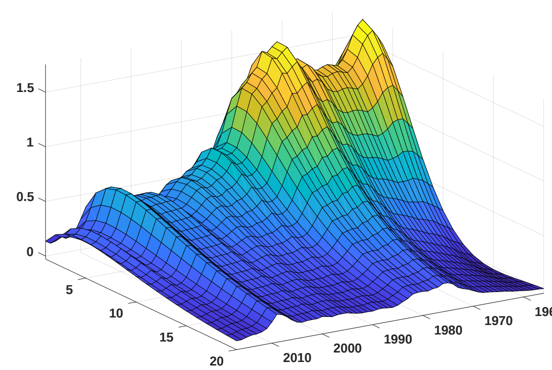

(f) Demand Shock: Output

Figure A.4: Dynamic Impulse Responses to Technology & Demand Shocks (LP, Hours & Output)

Note: Impulse Response Functions (IRF's) of per-capita hours, labour productivity and per-capita output from a 2-variable (viz., labour productivity growth and per-capita hours) time-varying long-run SVAR. Data is sourced from the BLS-LPC quarterly dataset for the U.S. business sector.

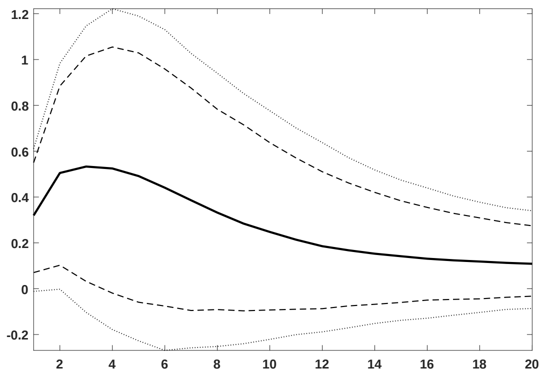

(a) Technology Shock: Hours

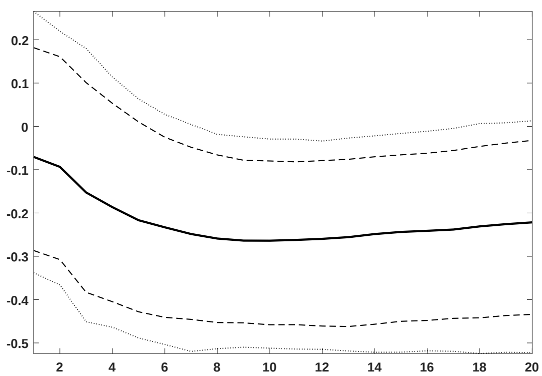

(b) Technology Shock: Labour Productivity

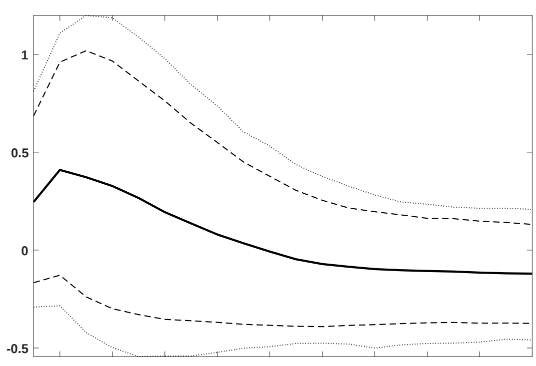

(c) Technology Shock: Output

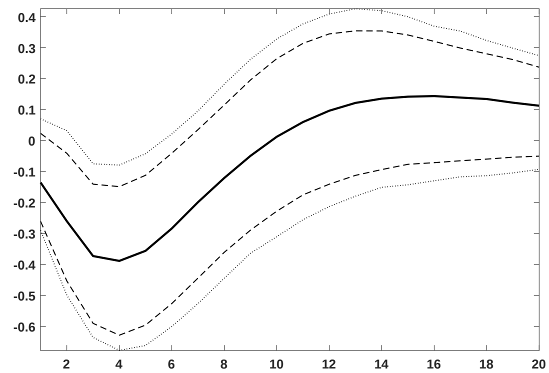

(d) Demand Shock: Hours

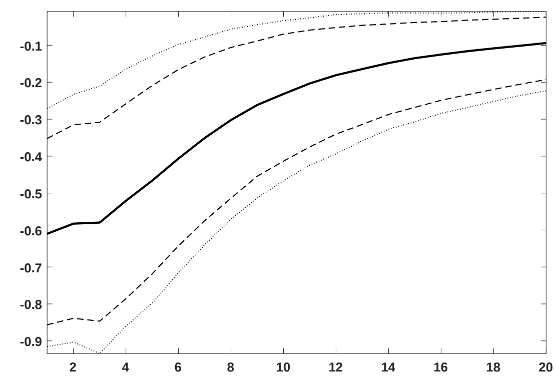

(e) Demand Shock: Labour Productivity

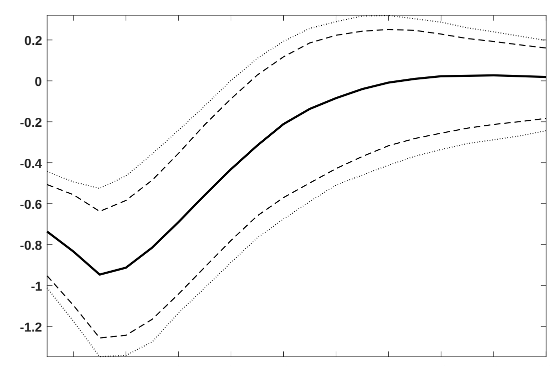

(f) Demand Shock: Output

Figure A.5: Difference in Impulse Responses between Pre- & Post-1984 (LP, Hours & Output)

Note: Post-1984 impulse response minus pre-1984 impulse response of per-capita hours, labour productivity and per-capita output from a 2-variable (viz., labour productivity growth and per-capita hours) time-varying long-run SVAR. The solid line is the difference in the impulse responses between pre- and post-1984 periods. The dotted and dashed lines are the 95% and 90% confidence intervals of the difference respectively. Data is sourced from the BLS-LPC quarterly dataset for the U.S. business sector.

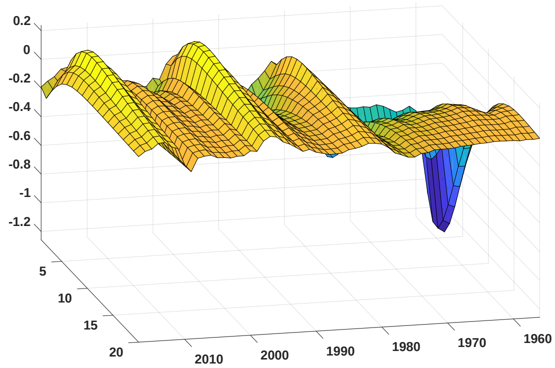

(a) Technology Shock: Hours

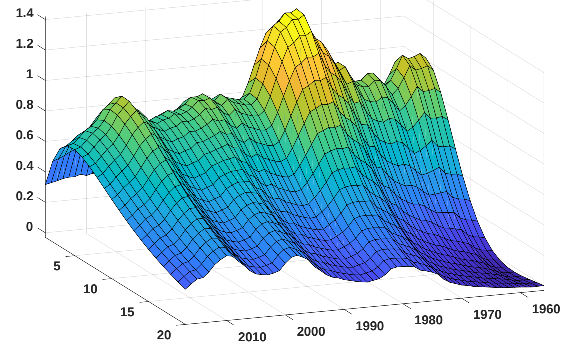

(c) Demand Shock: Hours

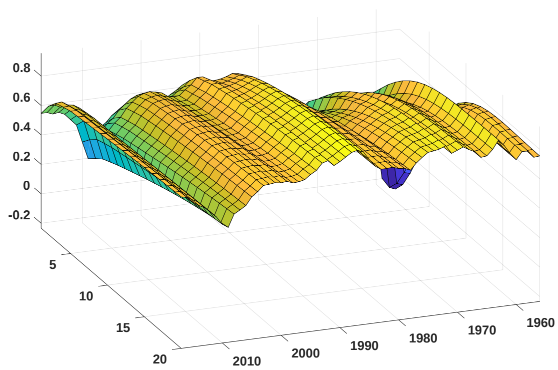

(b) Technology Shock: TFP

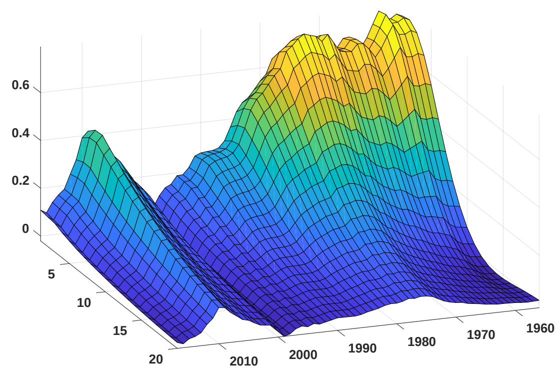

(d) Demand Shock: TFP

Figure A.6: Dynamic Impulse Responses to Technology & Demand Shocks (TFP & Hours)

Note: Impulse Response Functions (IRF's) of per-capita hours and TFP from a 2-variable (viz., TFP growth and per-capita hours) time-varying long-run SVAR. Hours data is sourced from the BLS-LPC quarterly dataset, TFP data is sourced from Fernald's quarterly TFP series, and quarterly civilian non-institutional population (16 years of age and older residing in the 51 U.S. states, who are not inmates of institutions, e.g., penal and mental facilities, homes for the aged, etc.) data is from the Employment Situation release of the BLS. All data correspond to the U.S. business sector.

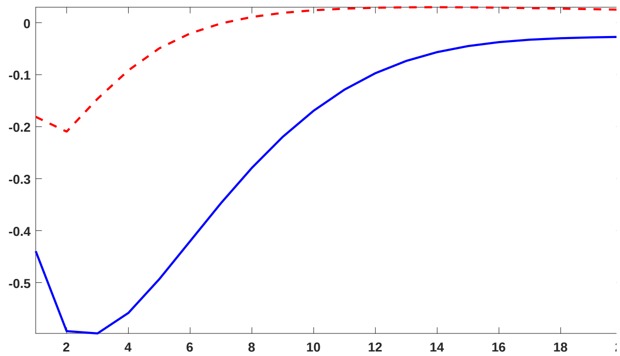

(a) Technology Shock: Hours

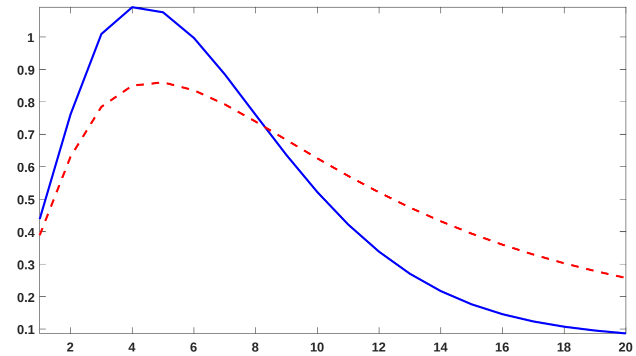

(c) Demand Shock: Hours

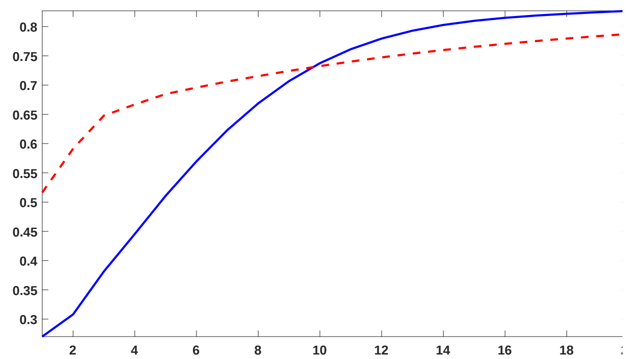

(b) Technology Shock: TFP

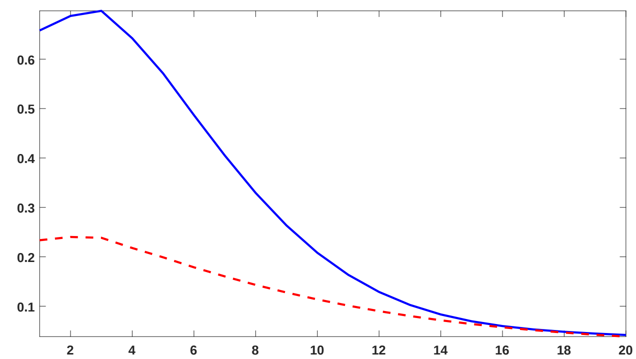

(d) Demand Shock: TFP

Figure A.7: Empirical Impulse Responses to Technology & Demand Shocks (TFP & Hours)

Note: Impluse Response Functions (IRF) for the pre-1984 period (1956-1983) are in blue, and the post-1984 (1984-2017) IRF's are in red dashed lines. Hours data is sourced from the BLS-LPC quarterly dataset, both types of TFP data are sourced from Fernald's quarterly series, and quarterly civilian non-institutional population (16 years of age and older residing in the 51 U.S. states, who are not inmates of institutions, e.g., penal and mental facilities, homes for the aged, etc.) data is from the Employment Situation release of the BLS. All data correspond to the U.S. business sector.

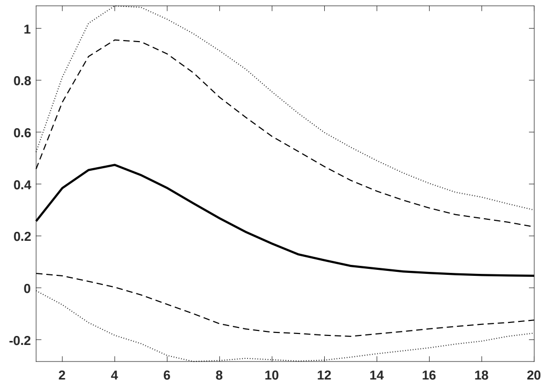

(a) Technology Shock: Hours

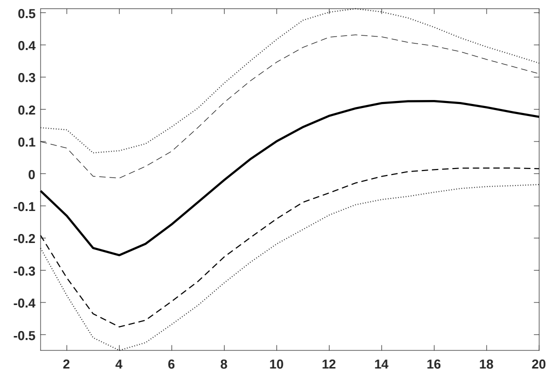

(c) Demand Shock: Hours

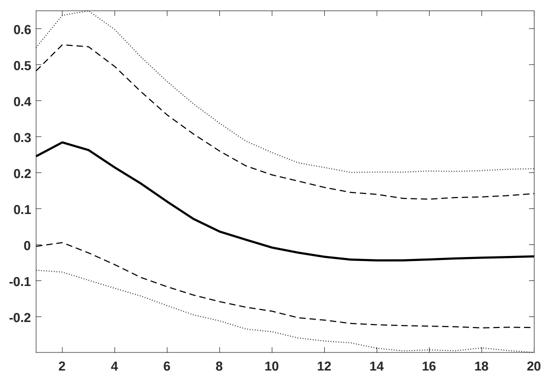

(b) Technology Shock: TFP

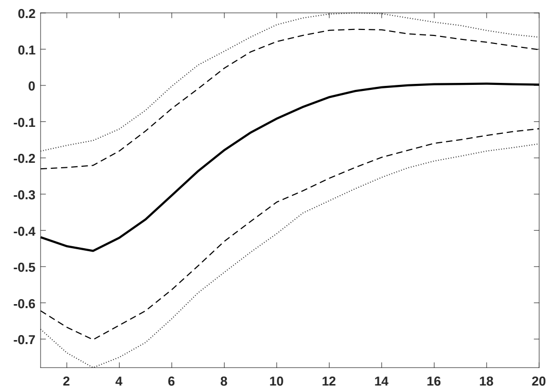

(d) Demand Shock: TFP

Figure A.8: Difference in Impulse Responses between Pre- & Post-1984 (TFP & Hours)

Note: Post-1984 impulse response minus pre-1984 impulse response of per-capita hours and TFP from a 2-variable (viz., TFP growth and per-capita hours) time-varying long-run SVAR. The solid line is the difference in the impulse responses between pre- and post-1984 periods. The dotted and dashed lines are the 95% and 90% confidence intervals of the difference respectively. Hours data is sourced from the BLS-LPC quarterly dataset, TFP data is sourced from Fernald's quarterly TFP series, and quarterly civilian non-institutional population (16 years of age and older residing in the 51 U.S. states, who are not inmates of institutions, e.g., penal and mental facilities, homes for the aged, etc.) data is from the Employment Situation release of the BLS. All data correspond to the U.S. business sector.

Appendix B

Appendix to Chapter 3

B.1 System of Log-Linearized Equations

Log-linearizing the model around a zero-inflation ($\bar{\pi^p} = 0$) steady-state with unit effort $\left(\bar{E} = 1\right)$ and employment rate, $\bar{N} = 0.62$, I get the following equations in log-deviation form, where the notation $\hat{x}_t$ is used to denote the deviation of logarithm of the variable X_t from its logged steady-state value $\bar{x}$.

$$\hat{y}_t = (1-\Theta)\,\hat{c}_t + \Theta\left(\hat{h}_t + \hat{g}_t\right) \tag{B.1}$$

$$\hat{y}_t = a_t + (1-\alpha)\left(\hat{n}_t + \psi\hat{e}_t\right) \tag{B.2}$$

$$\hat{n}_t = (1-\delta)\,\hat{n}_{t-1} + \delta\hat{h}_t \tag{B.3}$$

$$\hat{g}_t = \gamma\hat{h}_t \tag{B.4}$$

$$\hat{c}_t = \mathbb{E}_t\left(\hat{c}_{t+1}\right) - \hat{r}_t \tag{B.5}$$

$$\hat{r}_t = \hat{i}_t - \mathbb{E}_t\left(\pi^p_{t+1}\right) \tag{B.6}$$

$$\pi^p_t = \beta\mathbb{E}_t\left(\pi^p_{t+1}\right) - \lambda_p\hat{\mu}^p_t \tag{B.7}$$

$$\hat{\mu}^p_t = \left(\hat{y}_t - \hat{n}_t\right) - \left[(1-\Phi)\,\hat{\omega}_t + \Phi\hat{b}_t\right] \tag{B.8}$$

$$\hat{b}_t = \frac{1}{1-\beta(1-\delta)}\hat{g}_t - \frac{\beta(1-\delta)}{1-\beta(1-\delta)}\left[\mathbb{E}_t\left(\hat{g}_{t+1}\right) - \hat{r}_t\right] \tag{B.9}$$

$$\widehat{mrs}_t = \kappa\hat{c}_t + (1-\kappa)\left[\left(\hat{y}_t - \hat{n}_t - \hat{\mu}^p_t\right) + \frac{\iota}{1-\iota}\left(\hat{\omega}_t + \hat{n}_t - \hat{c}_t\right)\right] \tag{B.10}$$

$$\hat{\omega}_t = \hat{\omega}_{t-1} + \pi^w_t - \pi^p_t \tag{B.11}$$

$$\pi^w_t = \beta(1-\delta)\,\mathbb{E}_t\left(\pi^w_{t+1}\right) - \lambda_w\left(\hat{\omega}_t - \hat{\omega}^{target}_t\right) \tag{B.12}$$

$$\hat{\omega}^{target}_t = \Upsilon\widehat{mrs}_t + (1-\Upsilon)\left(\hat{y}_t - \hat{n}_t - \hat{\mu}^p_t\right) \tag{B.13}$$

$$\hat{i}_t = \rho\hat{i}_{t-1} + (1-\rho)\left(\phi_\pi\pi^p_t + \phi_y\hat{y}_t\right) + \phi_{\Delta y}\Delta\hat{y}_t + \nu_t \tag{B.14}$$

$$\hat{e}_t = \frac{1}{1+\phi}\left(\hat{y}_t - \hat{n}_t - \hat{\mu}^p_t - \hat{c}_t\right) \tag{B.15}$$

$$a_t = \rho_a a_{t-1} + \varepsilon^a_t \tag{B.16}$$

$$\nu_t = \rho_\nu\nu_{t-1} + \varepsilon^\nu_t \tag{B.17}$$

where $\Theta = \frac{\Gamma\left(\delta\bar{N}\right)^{1+\gamma}}{\bar{Y}}$, $\Phi = \frac{\bar{B}}{\bar{B}+\frac{\bar{W}}{\bar{P}}}$, $\kappa = \left(\frac{\chi}{1+\zeta}\right)\cdot\left(\frac{\bar{C}}{\overline{MRS}}\right)$, $\iota = \left(\frac{1+\phi}{1+\phi-\psi}\right)\cdot\left(\frac{\bar{W}\bar{N}}{\bar{P}\bar{C}}\right)$, $\Upsilon = \xi\left(\frac{\overline{MRS}}{\frac{\bar{W}}{\bar{P}}}\right)$, $\lambda_w = \frac{(1-\theta_w)(1-\beta\theta_w(1-\delta))}{\theta_w[1-(1-\Upsilon)(1-\Phi)]}$, and $\hat{\omega}_t = \hat{w}_t - \hat{p}_t$. The parameters ζ and χ are calibrated to satisfy unit effort in the steady-state ($\bar{E} = 1$) in a frictionless (no hiring cost) labour market. Furthermore, I take $\frac{\bar{W}\bar{N}}{\bar{P}\bar{C}} = \frac{1-\alpha}{1-\Theta}$.

B.2 Cyclical Moments of Capital and Factor Utilization

The model does not feature capital, rather includes only employment and effort. Since labour effort is not directly measurable in the data, one concern is that whatever is being labelled as 'effort' in the model is essentially capital, the missing factor of production. Therefore, it is important to distinguish between the business cycle dynamics of effort and capital. Using factor utilization rate as an empirically measurable proxy for effort, I show below how the cyclical moments of factor utilization in the data is qualitatively consistent with those of effort in the model, and they are different from those of capital.

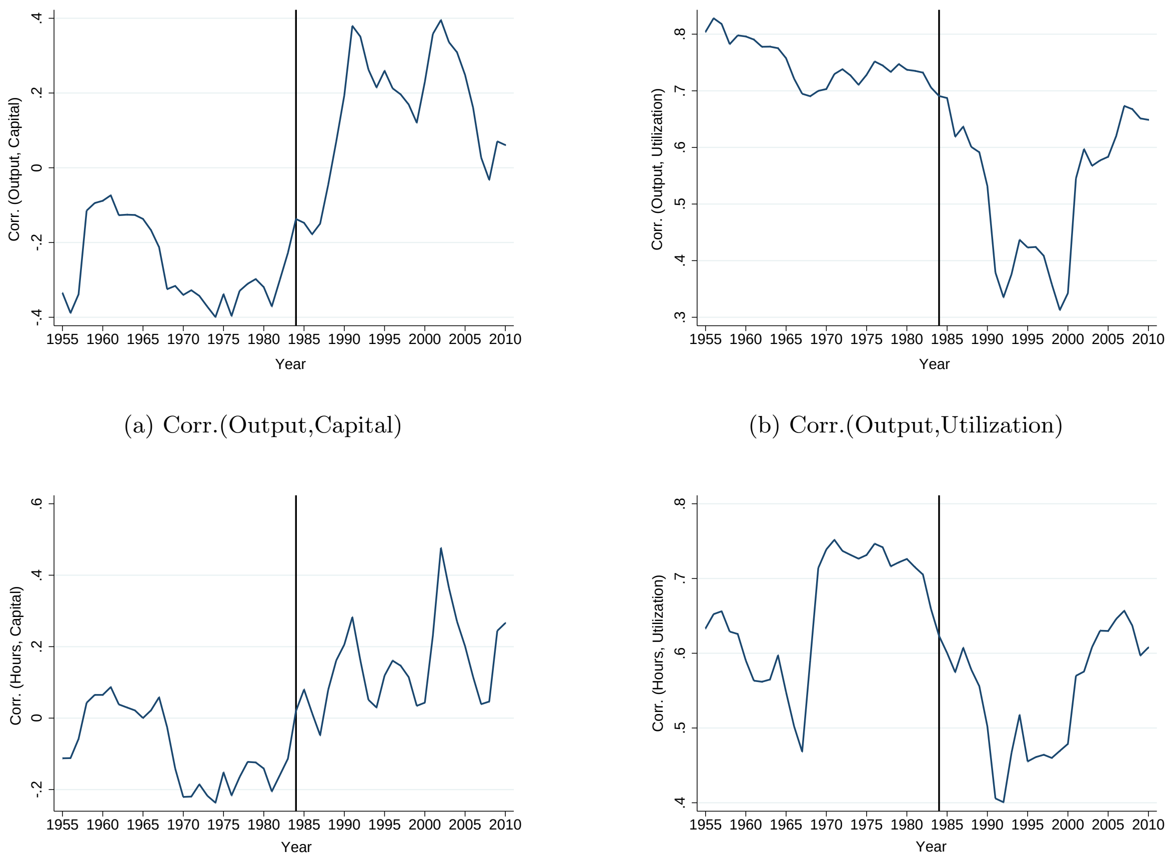

(a) Corr.(Output,Capital)

(b) Corr.(Output,Utilization)

(c) Corr.(Hours,Capital)

(d) Corr.(Hours,Utilization)

Figure B.1: Cyclical Correlations of Capital and Factor Utilization

Note: Data on quarterly growth rates of capital input, factor utilization, output and hours worked for the U.S. business sector are sourced from Fernald (2014). A centred rolling window of 15 years is used to calculate the second moments. Findings are robust to alternative choice of filters and window-sizes.

Looking at Panels (b) and (d) in Figure B.1, one can see that exactly around the time when productivity started losing its procyclicality, factor utilization also became more countercyclical. This fact was already presented in Table 2.1, where it was shown that the fall in aggregate TFP correlations with output and hours worked was driven by the reduced procyclicality of factor utilization and not utilization-adjusted TFP. However, it is immediately clear from the cyclical correlations of capital in Panels (a) and (c) of Figure B.1 that capital became more procyclical around the mid-1980s unlike factor utilization. Now, if the model implied correlations of effort with output and employment matches with those of factor utilization in the data then it can be argued that the role played by effort in the model is not the same as that of capital. Under the baseline calibration of the model (corresponding to column (4) of Table 3.3), correlation of effort with labour productivity fell by 0.37, which is qualitatively similar to that of factor utilization.

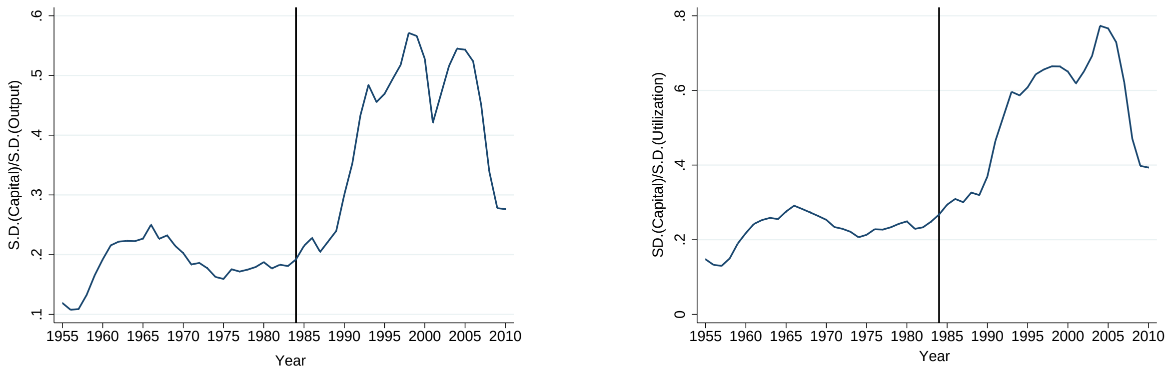

$\frac{\text{s.d.(Capital)}}{\text{s.d.(Output)}}$ $\frac{\text{s.d.(Capital)}}{\text{s.d.(Utilization)}}$

Figure B.2: Relative Volatility of Capital over the Business Cycle (1954-2010)

Note: Data on quarterly growth rates of capital input, factor utilization and output for the U.S. business sector are sourced from Fernald (2014). A centred rolling window of 15 years is used to calculate the second moments. Findings are robust to alternative choice of filters and window-sizes.

The volatility of capital relative to that of output and factor utilization rises sharply since the mid-1980s. It has already been shown that the relative volatility of employment has similarly rose. This further shows that the reliance on extensive margin of factor adjustment, for both labour and capital, has increased relative to the intensive margin of factor utilization. The model also predicts a substantial increase in the relative volatility of employment with respect to effort. All this evidence shows that the role of effort in the model is different from that of capital.

B.3 Volatility of Monetary Policy Shock

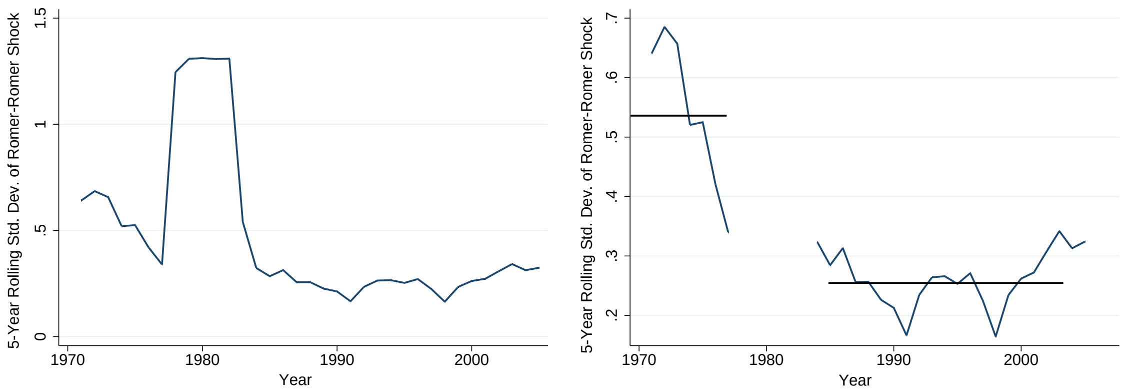

1971-2005 | 1971-1977, 1984-2005

Figure B.3: 5-Year Rolling Standard Deviation of Romer-Romer Monetary Shock

Note: Ignoring the sudden jump in volatility in the monetary policy shock between 1977 and 1982 as seen in Panel (a), the average standard deviation in the 1984-2005 period is roughly half of the average standard deviation during 1971-1977, as shown in Panel (b). Data for the Romer and Romer (2004) type monetary shock is sourced from Ramey (2016).

B.4 Data on Intellectual Property Products

I use the current-cost net capital stock of private non-residential fixed assets published by the Bureau of Economic Analysis (BEA) at the industry-level from 1947 through 2016. The data is disaggregated by asset type according to the classification by the National Income and Product Accounts (NIPA) — there are three major categories, namely, (i) equipment, with 39 sub-types, (ii) structures, with 32 sub-types, and (iii) intellectual property products (IPP), with 25 sub-types. The BEA typically does not include detailed estimates of different types of capital assets by industry in the tables published in the Survey of Current Business or the Fixed Assets and Consumer Durables volume because their quality is significantly lower than that of the higher level aggregates in which they are included. Compared to these aggregates, the detailed estimates are more likely to be either based on judgemental trends, on trends in the higher level aggregate, or on less reliable source data. Keeping this issue of data quality in mind, I will only use the share of aggregate IPP in total asset stock at the level of 24 U.S. industries. Below I present the time trend of the share of IPP in the total non-residential capital stock at current prices for the aggregate U.S. economy.

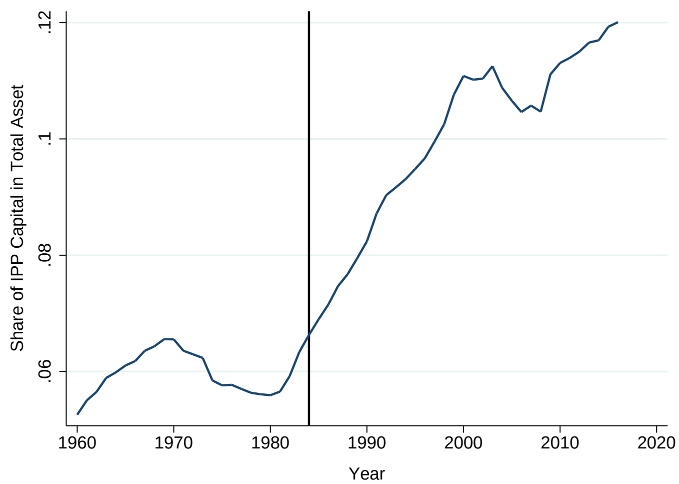

Figure B.4: Share of IPP in Total Non-Residential Capital Stock in the U.S. (1960-2016)

In order to give a clearer picture of what are the assets included under IPP, I provide below the complete list of NIPA asset-types that are categorized under IPP capital —
A. *Software*: Prepackaged, custom, and own account software
B. *Research & Development*: Pharmaceutical and medicine, other chemicals, semiconductor and other components, computers and peripheral equipment, communications equipment, navigational and other instruments, other computer and electronics, motor vehicles and parts, aerospace products and parts, and other manufacturing, scientific R&D services, software publishers, financial and real estate services, computer systems design and related services, all other non-manufacturing, private universities and colleges, and other non-profit institutions.
C. *Artistic Originals*: Theatrical movies, long-lived television programs, books, music, and other entertainment originals.

B.5 Relative Importance of Sector-Specific Shocks

Model:
$X_{i,t} = \lambda_i F_t + \varepsilon_{i,t}$
$X_{i,t}$: Observed growth rate of value added output or labour input for sector i at time t
F_t: Principal component of sectoral growth rates common to all sectors at time t
$\varepsilon_{i,t}$: Sector-specific growth rate for sector i at time t
Estimation:
Variance-covariance matrix of $X_{i,t}$, $V \equiv \Gamma\Lambda\Gamma'$ (Eigenvalue-Eigenvector Decomposition)
$F_t = X_{i,t}\Gamma_1$, where Γ_1 is the first eigenvector in Γ whose columns are sorted according to the ordering of the eigenvalues in Λ. The variance of F_t is interpreted as the aggregate economy-wide volatility (indicated as 'Common' in Tables B.1 and B.2), while that of $\varepsilon_{i,t}$ is the 'Sectoral' variance.

Table B.1: Components of Variance of Value Added Output Growth

	Pre-1983		**Post-1984**	
Dataset	Common	Sectoral	Common	Sectoral
BEA: 13 Sectors	92.93%	7.07%	68.30%	31.70%
KLEMS: 10 Sectors	48.14%	51.86%	4.42%	95.58%
KLEMS: 31 Sectors	17.96%	82.04%	5.15%	94.85%
IIP: 8 Sectors	94.98%	5.02%	87.21%	12.79%
IIP: 12 Sectors	70.89%	29.11%	31.49%	68.61%
IIP: 20 Sectors	30.63%	69.37%	42.18%	57.82%

Table B.2: Components of Variance of Labour Input Growth

	Pre-1983		**Post-1984**	
Dataset	Common	Sectoral	Common	Sectoral
CES: 14 Sectors	68.64%	31.36%	44.85%	55.15%
BEA: 13 Sectors	92.31%	7.69%	74.61%	25.39%
KLEMS: 10 Sectors	78.28%	21.72%	50.87%	49.13%
KLEMS: 31 Sectors	89.36%	10.64%	91.14%	8.86%

It should be noted that most of the rise in the relative importance of sector-specific variance was due to the drop in the variance of the common component, while the sectoral component remained largely constant between the pre and post-1984 periods. In other words, one can conclude that the drop in variance of output and labour inputs during the Great Moderation was mostly due to falling volatility of the aggregate shocks rather than sectoral ones. However, this is not true when a 31-industry-split is considered — both the common and sectoral variances decline in that case.

Appendix C

Appendix to Chapter 4

C.1 Derivation of the Consumption Process

In this appendix we derive the analytical approximation of the optimal consumption processes. Assuming a quadratic utility function and $\beta(1+r)=1$, we solve the maximization problem (4.5) and derive consumption at time t as the annuity value of lifetime resources, as follows:

$$C_{f,t}=\frac{r}{(1+r)-(1+r)^{-(T-t)}}\left[A_{f,t}+\sum_{j=0}^{T-t}\left(\frac{1}{1+r}\right)^{j}\mathbb{E}_t\left(E_{f,t+j}\right)+\sum_{j=0}^{T-t}\left(\frac{1}{1+r}\right)^{j}\mathbb{E}_t\left(N_{f,t+j}\right)\right]$$

To express consumption expenditure in terms of logs, we use a first order Taylor series approximation of the logarithm of each variable around unity. For any variable x, $ln(x)\simeq ln(1)+\frac{x-1}{1}=x-1$.[89] Using this approximate relationship $x\simeq 1+ln(x)$, and denoting $ln\left(C_{f,t}\right)$, $ln\left(A_{f,t}\right)$, $ln\left(E_{f,t}\right)$ and $ln\left(N_{f,t}\right)$ by $c_{f,t}$, $a_{f,t}$, $e_{f,t}$ and $n_{f,t}$ respectively, and using the time-series processes we assumed for $a_{f,t}$, $e_{f,t}$ and $n_{f,t}$, we get,

$$\begin{aligned}
1+c_{f,t} \quad &\simeq \quad (1+\bar{e}_f)+(1+\bar{n}_f)+\\
\tfrac{r}{(1+r)-(1+r)^{-(T-t)}} \quad &\left\{(1+a_{f,t})+\left[\sum_{j=0}^{T-t}\left(\frac{1}{1+r}\right)^{j}\mathbb{E}_t\left(\zeta_{f,t+j}\right)+\sum_{j=0}^{T-t}\left(\frac{1}{1+r}\right)^{j}\mathbb{E}_t\left(u_{f,t+j}\right)\right]\right\}\\
\Longrightarrow c_{f,t} \quad &\simeq \quad 1+\frac{r}{(1+r)-(1+r)^{-(T-t)}}\left[(1+a_{f,t})+(\zeta_{f,t}+u_{f,t})\right]+\bar{e}_f+\bar{n}_f
\end{aligned}$$

Let $q_{f,t}\equiv 1+\alpha_t(r)\left(1+a_{f,t}\right)$, with $\alpha_t(r)=\frac{r}{(1+r)-(1+r)^{-(T-t)}}$. Then we can write the approximate log-consumption processes for an individual as:

$$c_{f,t}=q_{f,t}+\bar{e}_f+\bar{n}_f+\alpha_t(r)\left(\zeta_{f,t}+u_{f,t}\right)$$

For a large enough T relative to t, $\alpha_t(r)$ can be approximated by $\alpha(r)=\frac{r}{1+r}$. Thus, for individuals who are sufficiently away from their demise, we can approximate their log-consumption

[89] This approximation holds only for values of x close to unity. Since in the empirical implementation of the model, we de-mean all the log variables ($\ln x$), this approximation is valid.

as:

$$c_{f,t} = q_{f,t} + \bar{e}_f + \bar{n}_f + \alpha\left(r\right)\left(\zeta_{f,t} + u_{f,t}\right) \tag{C.1}$$

C.1.1 CRRA Utility Function

Relaxing the assumption of a quadratic utility function, we can still arrive at the same log-consumption equation as (C.1) with a more general utility function, after a linear approximation of the Euler equation. For example, in the case of constant relative risk aversion (CRRA) utility function, the Euler equation is given by $C_{f,t}^{-\tau} = \beta\left(1+r\right)\mathbb{E}_t\left(C_{f,t+1}^{-\tau}\right)$, where $\tau > 0$ is the parameter capturing the degree of risk aversion as also the intertemporal elasticity of substitution. Maintaining the assumption $\beta\left(1+r\right) = 1$, we get from the Euler equation $\mathbb{E}_t\left[\left(\frac{C_{f,t+1}}{C_{f,t}}\right)^{-\tau}\right] = 1$. We define the function $h\left(g_c\right) = \left(1+g_c\right)^{-\tau}$, where $g_c = \frac{C_{f,t+1}}{C_{f,t}} - 1$ such that $\mathbb{E}_t\left[h\left(g_c\right)\right] = 1$. A first order Taylor series expansion of $h\left(g_c\right)$ around $g_c^* = 0$ yields $h\left(g_c\right) \approx 1 - \tau g_c$. Taking expectations on both sides of this approximate equation, we get $\mathbb{E}_t\left(g_c\right) = 0$, implying $C_{f,t} = \mathbb{E}_t\left(C_{f,t+1}\right)$. This is exactly the same as the Euler equation that one obtains from quadratic utility function without any approximation. Now, since we did not derive explicitly the consumption expression from this Euler equation in the paper, we provide the derivation here. Iterating forward the per-period budget constraint $A_{f,t+1} = \left(1+r\right)\left(A_{f,t} + Y_{f,t} - C_{f,t}\right)$ (where $Y_{f,t} = E_{f,t} + N_{f,t}$) by one period and combining it with the Euler equation $C_{f,t} = \mathbb{E}_t\left(C_{f,t+1}\right)$, we get,

$$\begin{aligned}
\left(1+\frac{1}{1+r}\right)C_{f,t} &= A_{f,t} - \left(\frac{1}{1+r}\right)^2 \mathbb{E}_t\left(A_{f,t+2}\right) + \left[Y_{f,t} + \frac{1}{1+r}\mathbb{E}_t\left(Y_{f,t+1}\right)\right] \\
&\vdots \\
\Rightarrow \left[1+\frac{1}{1+r}+\left(\frac{1}{1+r}\right)^2 + ...\infty\right]C_{f,t} &= A_{f,t} - \lim_{k\to\infty}\left(\frac{1}{1+r}\right)^k \mathbb{E}_t\left(A_{f,t+k}\right) + \sum_{j=0}^{\infty}\left(\frac{1}{1+r}\right)^j \mathbb{E}_t\left(Y_{f,t+j}\right) \\
\Longrightarrow \left[\frac{1+r}{r}\right]C_{f,t} &= A_{f,t} + \sum_{j=0}^{\infty}\left(\frac{1}{1+r}\right)^j \mathbb{E}_t\left(Y_{f,t+j}\right) \\
\Longrightarrow C_{f,t} &= \frac{r}{1+r}\left[A_{f,t} + \sum_{j=0}^{\infty}\left(\frac{1}{1+r}\right)^j \mathbb{E}_t\left(Y_{f,t+j}\right)\right]
\end{aligned}$$

Note that in the above derivation we have assumed the No-Ponzi condition that prevents an individual from continuously borrowing and rolling over his debt to future periods, $\lim_{k\to\infty}\left(\frac{1}{1+r}\right)^k \mathbb{E}_t\left(A_{f,t+k}\right) = 0$.

C.2 Data and Sampling

The Panel Study of Income Dynamics (PSID) is administered by the University of Michigan's Survey Research Center (SRC). This longitudinal survey began in 1968 with a national probability sample of almost 5,000 U.S. families. The sampled families were re-interviewed annually between 1968 and 1997. After 1997 they were re-interviewed biennially. We focus our study only on the non-Latino, non-immigrant households within the SRC component of the PSID, and exclude those in the Survey of Economic Opportunity (SEO) component where poor households were over-sampled.

PSID data have been used by different authors for intergenerational analyses because, by design, this survey follows the children of original sample members when they become independent from their original family. This allows to follow children from the original sample as they grow into adulthood and become household heads themselves. To reduce noise due to weak labour market participation and marital status, our main analysis for household heads focuses on observations for married male individuals between 25 and 65 years of age, who have at least 5 years of data in the PSID, have non-negative labour earnings and total family income, work for less than 5840 hours annually, have wages greater than half of the federal minimum wage, and do not have annual earnings growth rates of more than 400 percent. Our analysis pertains to children born between 1952 and 1981. To avoid over-representation of children who left their homes at a later stage of their lives, the sample excludes children born before 1952 (that is, those children who were older than 16 at the time of the first 1968 PSID interview). The first year in which child income is observed is 1977 (as reported in the 1978 interview) - the year in which the 1952 birth-cohort reached age 25. Consequently, we can observe the 1952 cohort between ages 25 and 62, while the 1981 cohort can only be observed between ages 25 and 33 years. Parents who are older than 65 are dropped from the analysis to avoid complications related to retirement decisions. In robustness checks, we consider various alternative samples, e.g., restrict age range from 30 to 40 years for both parents and children, and look at different cohorts of children separately. Our model estimates remain qualitatively similar under all these alternative samples.

The labour earnings data for the male household head and his wife, and the total transfer income data for the couple are readily available for most survey rounds of the PSID. In contrast, the family consumption data is quite sparse across the survey years and not presented as a single variable in the PSID. Different consumption expenditure categories have to be suitably summed up (using appropriate weights depending on the frequency of consumption in a particular category, e.g., yearly, monthly, weekly, etc.) to arrive at an aggregate measure of consumption expenditure.

There are 11 major categories of consumption variables, namely, (i) food, (ii) housing, (iii) child-care, (iv) education, (v) transportation, (vi) healthcare, (vii) recreation and entertainment, (viii) trips and vacation, (ix) clothing and apparel, (x) home repairs and maintenance, and (xi) household furnishings and equipment. Of these, food and housing are most consistently observed across the years - expenditure on food is observed from the 1968 interview through the 2015 interview, barring only 1973, 1988 and 1989. Housing expenditure is observed in all years except 1978, 1988 and 1989. Child-care expenditure data is available for 25 rounds of interview - 1970-1972 (3 interview years), 1976, 1977, 1979 and 1988-2015 (19 interview years). Education, transportation and health-care are only reported by the last 9 PSID interviews (biennially from 1999 through 2015). The rest of the categories from (vii) through (xi) are observed for only the last 6 interviews (biennially from 2005 to 2015).

The uneven availability of expenditure categories in different waves of the PSID suggests that a simple sum of the expenditure categories for different years would not provide an accurate approximation of total consumption because every year reports different subsets of consumption expenditures. There are two ways to account for this problem in the calculation of the total consumption variable: either take the measure of consumption to be equal to just the expenditure on food, the most consistently observed category (although that would ignore variation in the consumption of non-durable goods other than food); or impute the consumption of the missing categories.

C.2.1 Imputation of Consumption Expenditure Data

To assess the quality of consumption survey data, Andreski, Li, Samancioglu, and Schoeni (2014) compare expenditure data from the Consumption Expenditure Survey (CEX) and the PSID. They find that expenditures in individual categories of consumption may vary non-trivially across the two datasets, e.g., reported home repairs and maintenance expenditures are approximately twice as large in the PSID as the are in the CEX, and the PSID home insurance expenditures are 40 to 50 percent higher than their CEX counterparts. However, despite these inconsistencies within individual categories (due to differences in survey methodologies and sampling techniques), Li, Schoeni, Danziger, and Charles (2010) show that the average expenditure since 1999 in PSID and CEX have been fairly close to each other. Moreover, the consumption expenditures in the two datasets vary in a similar way with observable household characteristics like age of household head, household size, educational attainment, marital status, race and home ownership. This average consistency between PSID and CEX data, as well as the fact that total consumption seems to be close to the aggregate consumption estimates in the NIPA (National Income and Product Accounts) data, suggests that PSID expenditure data can be used to draw information about households consumption behavior.

Attanasio and Pistaferri (2014) (henceforth AP) suggest to impute consumption data for the missing consumption categories in the PSID before 1999 by using the more detailed data available post-1999. Their backward extrapolation is consistent with theories of consumer demand in the sense that the allocation of total resources spent in a given period over different commodities is made dependent on relative prices and taste-shifters, e.g., demographic and socio-economic variables. However, this specification implicitly assumes homotheticity of consumer preferences over different commodities. To relax that assumption, we include log total income in the imputation regression as a control. We use this slightly modified approximated demand system to total consumption expenditures before 1999:

$$ln\,(N_{ft}) = Z'_{ft}\omega + p'_t\pi + g(F_{ft};\lambda) + u_{ft}, \tag{C.2}$$

where N is consumption net of food expenditure, Z are the socioeconomic controls (viz., dummies for age, education, marital status, race, state of residence, employment status, self-employment, head's hours worked, homeownership, disability, family size, and the number of children in the household) and total family income, p are the relative prices (the overall CPI and the CPIs for food at home, food away from home, and rent), F is the total food expenditure (i.e., sum of food at home, food away from home, and food stamps) that is observed in the PSID consistently through the years, $g(.)$ is a polynomial function, and u is the error term. The subscripts f and t denotes family identity and year respectively. This equation is estimated using data from the 1999-2015 PSID waves, where the net consumption measure N_{ft} is the sum of annualized expenditures on home insurance, electricity, heating, water, other miscellaneous utilities, car insurance, car repairs, gasoline, parking, bus fares, taxi fares, other transportation, school tuition, other school expenses, child care, health insurance, out-of-pocket health, and rent. While performing the imputation we skip the consumption expenditure categories that were added to the PSID from the 2005 wave. This is done to keep the measure of consumption consistent over the years and to also maximize the number of categories that can be used. Moreover, the categories added from the 2005 wave collectively constitute a very small fraction of total consumption. In the definition of net consumption we have excluded food expenditure to avoid endogeneity issues in the regression. The measure for rent equals the actual annual rent payments for renters and is imputed to 6% of the self-reported house value (Flavin and Yamashita (2002)) for the homeowners.

After estimating the logarithm of the net consumption equation by running a pooled OLS regression on equation (C.2), we construct a measure of imputed total consumption as follows

$$\hat{C}_{ft} = F_{ft} + exp\left\{Z'_{ft}\hat{\omega} + p'_t\hat{\pi} + g\left(F_{ft};\hat{\lambda}\right)\right\}. \tag{C.3}$$

This measure is corrected for inflation by dividing it by the overall CPI. Finally the measure is transformed into adult-equivalent values using the OECD scale, $(1 + 0.7(A - 1) + 0.5K)$, where A is the number of adults and K the number of children in the household unit.

A key question is how well the imputed consumption values match with the observed values during the period when both data series are available. A natural choice for a measure of the goodness of fit is the R^2 of the regression (C.2), which is found to be 0.47. However, what we are really interested in is matching the standard deviations of the observed and imputed series because we would be using only the second order moments of income and consumption for estimating our model in Section 4.2. Like AP, we find that our imputed consumption series can match the observed series quite closely in terms of standard deviation, and similarly well for a more general non-linear measure like the Gini coefficient. Figure C.1 presents the Gini coefficients (normalized to their initial values in 2006) of the logs of imputed and actual consumption (in Panel C), and also compares the standard deviations of actual and imputed consumption with those of real income and labour earnings (in Panels A, B, and D). The top-coded values for total family income and the household heads' labour earnings in the PSID are replaced with the estimates obtained from fitting a Pareto distribution to the upper tail of the corresponding distribution.

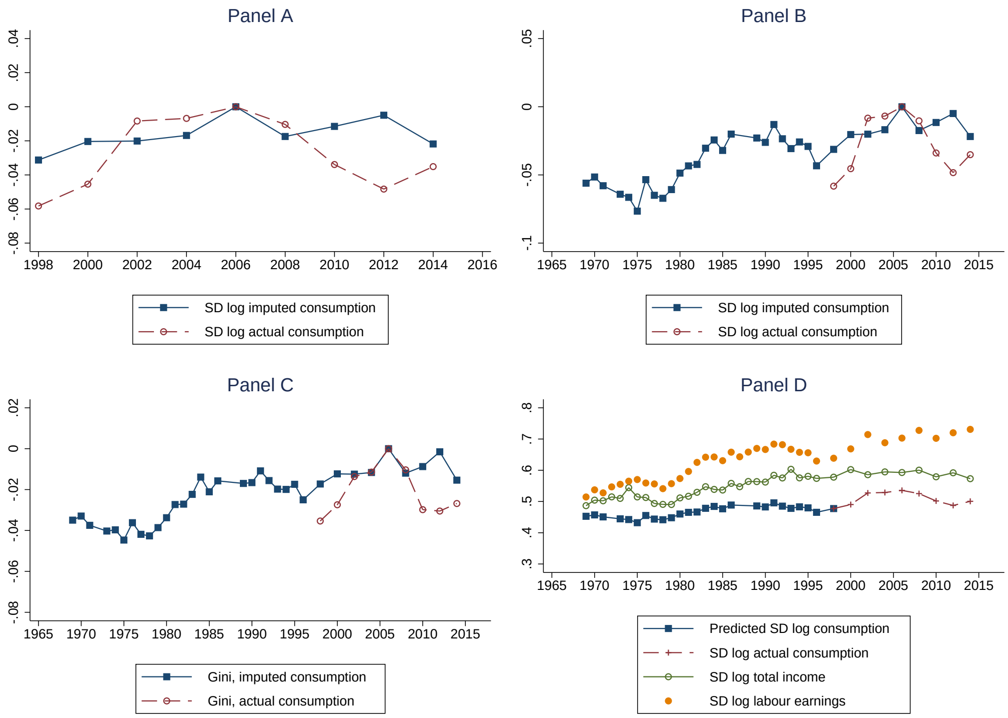

Figure C.1: Quality Assessment of Consumption Imputation
Note: In Panels A, B and C, series are normalized to values in 2006 for ease of comparison.

C.3 Intergenerational Persistence: Reduced-Form Evidence

In this appendix, we present some reduced form evidence of the time trends and cross-sectional heterogeneity of intergenerational persistence in earnings, as common in the literature, and also consumption, which is more closely tied to welfare.

C.3.1 The Evolution of Intergenerational Elasticities

A natural way to measure the impact of parental economic circumstances on a child's adult outcomes is to estimate the intergenerational elasticity of such outcomes. By definition, this elasticity measures the percentage change in the child's variable following one percentage change in the corresponding parental variable, and is obtained by regressing a logged measure of the child's variable on its parental counterpart.

We are interested in knowing the persistence in permanent earnings and consumption, but we do not directly observe the long-term (permanent) earnings and consumption of any individual. An adult child's earnings are observed only over a limited range of ages. Hence we must proxy these life-cycle variables by some function of the current (yearly) variables that are actually observable.[90] As

[90] A simpler way of dealing with this issue is to take into account the relevant variable at a particular age (say 30)

in Lee and Solon (2009) we use adult children's data for all the available years, along with a full set of age controls. We centre the child's age around 40 years to minimise the bias from heterogeneity in growth rates, and interpret the estimated intergenerational elasticity as an average value as successive cohorts of children pass through age 40.[91] In fact, these intergenerational elasticities at age 40 (for a given year) can be interpreted as an asymmetrical moving average of the cohort-specific elasticities for the cohorts of adult children who are observed for that particular year. It is asymmetrical because the older cohorts weigh more in a particular year's estimate owing to the fact that cohorts enter as they turn 25 years of age but never leave till the end of the PSID dataset.[92]

We also need to use a suitable proxy for the long-run parental variable serving as the principal regressor. Using the current measure of the parental variable would introduce an attenuation bias in the estimation of the long-term intergenerational elasticity of the child's variable. As in Lee and Solon (2009), we use the average log annual value of the parental variable over the years when the child was between age 15 and 17 as a proxy for the long-run value of the parent's process. We choose 15 years as the starting child age for a parental observation because our focus is on how parental circumstances in the formative years affects outcomes.[93] An alternative would have been to take the average of the parental variable (earnings or consumption) for the parents' entire lifetime (till 65 years of age). This would confound a number of effects, in particular, the effect of parental outcomes when children are at home with realisations of parental outcomes after children left home. The latter contemporaneous pass-through may be important for consumption smoothing across generations, but conceptually it is a different mechanism. A further issue with using the average over the entire lifetime is that this would impose that siblings born at different life-stages of the parent face the same parental inputs. Obviously, the age of the parents of different children born in a particular cohort will not be the same when the children reach the age range between 15 and 17. Therefore, we also control for the age of the parental household head.

We define the dependent variable ζ_{fht} as the outcome variable — earnings or consumption, of the child f born in year h observed in year t. We run the regression:

$$\zeta_{fht} = \mu D_t + \beta_t x_{fh} + \gamma a^p_{fh} + \delta a^k_{fht} + \epsilon_{fht} \tag{C.4}$$

The regressor, x_{fh} is the average value of the parent's outcome variable when the child f from cohort h is between 15 and 17 years of age. As controls, we include year dummies D_t, and quartics

for all children. This approach is adopted, for example, by Mayer and Lopoo (2005). The downside of conditioning on a specific age is that one has to throw out much valuable information (that is, all the data available for other ages). Moreover, transitory shocks occurring at the specific age may introduce some bias in the estimated parameter.

[91] Classical measurement error in the dependent variable (here, the child variable) is usually not a problem. However, Haider and Solon (2006) shows that using current variables as a proxy for a child's permanent (lifetime) earnings or income may entail non-classical measurement error but the extent of the measurement error bias in the left-hand-side variable is the lowest if the current variable is measured at around age 40. So, we centre the child's age around age 40.

[92] This asymmetry can be easily removed by making cohorts exit after a certain age, but that would lead to missing out on valuable information for those omitted cohorts. An alternative to this time-conditional estimation is to estimate cohort-specific elasticities using lifetime average of earnings (or consumption) for the adult children.

[93] Data availability then implies that is the oldest cohort of children are those born in 1952, with available parental observations starting from 1967 (documented in the 1968 interview).

in the average parental age when the child is age 15-17 years, a^p_{fh}, and also quartics in the age of the child in year t, centred around 40 years (that is, a quartic in $t-h-40$), a^k_{fht}. The error term ϵ_{fht} reflects factors like luck in labour and marriage markets, intergenerational transmission of genetic traits and other environmental factors (see Peters (1992)). We allow the coefficient β to vary by year to capture the time variation in intergenerational persistence. It should be noted that the choice of the normalization age for a^k_{fht} affects the point estimate of β_t in each year but not the time trend. In Table C.1 we report the actual year-specific estimates from 1990 through 2010. We can obtain estimates starting from 1977 onwards, but in earlier years of the PSID the average age of the children samples is quite low, as we only observe independent children for very few years. This is problematic because one would have to rely on extremely short snapshots of early adulthood to infer child outcomes. For this reason we only report point estimates of the elasticities from the year 1990 onwards. This guarantees that the cross-section of children in any given year includes a larger number of individuals at later stages of their working life. This also guarantees that children panels are longer, and hence less susceptible to initial conditions bias. It is interesting to note that the estimated elasticities lie in a fairly narrow range in the last 30 years. This absence of either a positive or a negative trend is the basis of our time-stationary model of economic persistence in Section 4.2.

Table C.1: Estimates of Intergenerational Elasticities by Year

Year	**Head Earnings**	**Total Consumption**	**Food Consumption**
1990	0.30***	0.48***	0.25***
1991	0.34***	0.45***	0.24***
1992	0.29***	0.47***	0.27***
1993	0.30***	0.48***	0.29***
1994	0.29***	0.49***	0.25***
1995	0.29***	0.48***	0.27***
1996	0.25***	0.45***	0.25***
1998	0.24***	0.44***	0.24***
2000	0.30***	0.45***	0.25***
2002	0.31***	0.48***	0.23***
2004	0.29***	0.41***	0.19***
2006	0.30***	0.46***	0.23***
2008	0.35***	0.47***	0.26***
2010	0.37***	0.49***	0.29***

Note: ***, ** and * denote statistical significance at 1%, 5% and 10% levels respectively. Standard errors (not reported) are clustered at the level of the unique parent identity.

C.3.2 Heterogeneity of Intergenerational Persistence

An alternative way to study the extent of intergenerational economic persistence is through *mobility matrices*. Mobility matrices show the heterogeneity in intergenerational persistence across the income or consumption distribution that is averaged out in the regression analysis above and the

GMM analysis later on. The basic idea is to study the probability that an adult child will fall into various quantiles in the income or consumption distribution, given the quantile in which the parent of that child belonged. If the probability of a child being placed in the same quartile as the parent is high, we say that intergenerational persistence is high for that quartile of the distribution. If there were to be perfect intergenerational mobility then each cell in the mobility matrix would have a conditional probability of 25 percent, and on the other hand if there were perfect persistence in intergenerational well-being then all the diagonal cells would read 100 percent while the off-diagonal cells would have a zero probability.

To accomplish the construction of such mobility matrices we first regress parental earnings (or consumption) on the full set of year dummies and the quartic of parental age. The residuals from these regressions are then averaged across the years for each parent and these average residuals are finally used to place each parent in one of the four quartiles of the parental distribution. Similar exercise with the adult children is performed, and finally the two quartile positions of the parents and children are cross-tabulated. A cell $c_{i,j}$ in a mobility matrix at the intersection of the i^{th} row and the j^{th} column $\forall i, j = 1(1)4$ is given by

$$c_{i,j} = Prob\,[child \in Q_{c,i} |\, parent \in Q_{p,j}] \times 100$$

where $Q_{c,i}$ denotes the i^{th} quartile of the child distribution and $Q_{p,j}$ denotes the j^{th} quartile of the parental distribution. One should note that the sum of each column in a mobility matrix must add up to 100. This is because the sum is essentially the integration of the conditional distribution for the child over the entire range of that distribution. However, the sum of each row need not add up to 100.

Mobility Matrix of Head Earnings

Child \ **Parent**	$Q_{p,1}$	$Q_{p,2}$	$Q_{p,3}$	$Q_{p,4}$
$Q_{c,1}$	**45.98**	27.88	17.29	9.56
$Q_{c,2}$	25.41	**29.64**	27.17	15.93
$Q_{c,3}$	19.75	24.80	**30.44**	23.10
$Q_{c,4}$	8.86	17.69	25.10	**51.41**

Mobility Matrix of Total Consumption

Child \ Parent	$Q_{p,1}$	$Q_{p,2}$	$Q_{p,3}$	$Q_{p,4}$
$Q_{c,1}$	**53.02**	27.79	9.75	4.95
$Q_{c,2}$	26.53	**32.04**	25.65	13.65
$Q_{c,3}$	16.28	26.51	**35.40**	23.55
$Q_{c,4}$	4.17	13.67	29.20	**57.84**

Mobility Matrix of Food Consumption

Child \ Parent	$Q_{p,1}$	$Q_{p,2}$	$Q_{p,3}$	$Q_{p,4}$
$Q_{c,1}$	**40.00**	26.24	21.53	10.17
$Q_{c,2}$	27.03	**30.19**	20.26	20.75
$Q_{c,3}$	21.11	24.00	**32.07**	23.30
$Q_{c,4}$	11.86	19.57	26.14	**45.78**

The mobility matrices for household head's labour earnings, total family consumption and food consumption are provided above. There are two important observations to be made from the tables. First, the mobility matrix of labour earnings show more mobility than that of total consumption. This implies the presence of other channels of intra-family linkages in consumption that are over and above earnings. Note that this finding is consistent with the intergenerational elasticities above. The contributions of these different channels of persistence will be explicitly quantified in the more structural model in Section 4.2. Secondly, there is a lot of heterogeneity in economic persistence across the conditional distributions, with the most persistence being observed at the two tails of the distributions, e.g., among children whose parents were in the lowest quartile of the parental distribution, at least about 39 percent are also in the lowest quartile. There is much more mobility in the middle of the distributions.

Mobility matrices, while good at highlighting distributional heterogeneity in intergenerational persistence, as such cannot provide a summary statistic for measuring the overall mobility in the economy. Using the fact that in the case of perfect persistence the mobility matrix is nothing but the identity matrix of size m, where m is the number of quantiles used to construct the mobility matrix (in our case of quartiles, $m = 4$), (Shorrocks, 1978) provides a simple measure of the distance of the estimated mobility matrix (M) from the identity matrix as follows:

$$\textbf{Normalized Trace Index}, NTI = \frac{m - trace(M)}{m-1}$$

The NTI measure is ***0.81*** *for the labour earnings transition matrix, while that for total consumption expenditure and food consumption stand lower at* ***0.74*** *and* ***0.84*** *respectively. This corroborates the higher persistence of total consumption than earnings and food consumption.*

C.4 Empirical Moments

The GMM minimizes the distance between the empirical moments and the analytical moments implied by the statistical model. If the parameters were exactly identified then the GMM estimates would be nothing but the solution of the system of moment restrictions. However, with over-identification, the GMM becomes relevant in the sense that it minimizes the error from all over-identifying restrictions. Hence, it is important that we study the empirical moments which essentially gives the estimates via the GMM. Below we present the cross-sectional empirical moments for the baseline case along with the internal fit of the model.

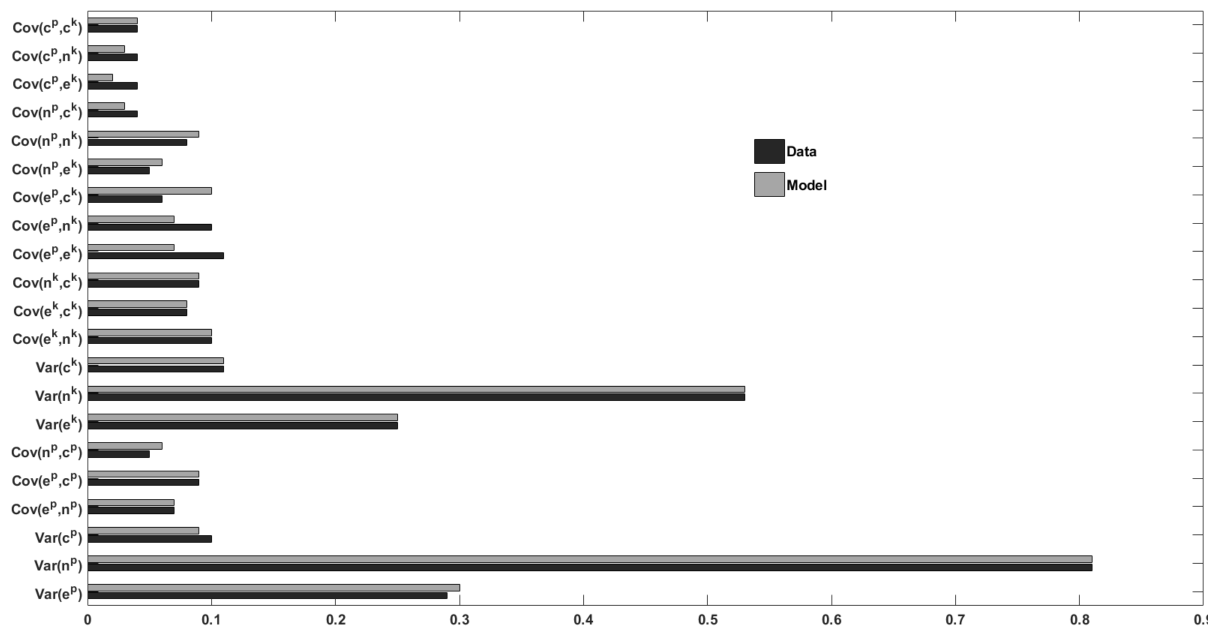

Figure C.2: Internal Fit of Baseline Model

Note: Both the data and the model estimates correspond to the baseline case where the raw data is purged of only birth-cohort and year fixed effects. The average age for parents is 47 years, while that for children is 37 years.

C.5 Supplementary Results

C.5.1 Role of Observable Characteristics in Intergenerational Persistence

How much of the intra-family linkages in earnings, other income and consumption can be explained by observable characteristics of the two generations? Observables like race and educational attainment has long been argued to be significant determinants of intergenerational mobility.

Table C.2: Persistence of Observable Characteristics

Observed Variable	**Persistence**
Family Size	0.32
State of Residence	0.71
No. of Children	0.38
Employment Status	0.86
Race	0.98
Education	0.50

Table C.2 shows the high degree of persistence in a host of observable characteristics across the two generations in our sample. So a natural question to ask is — if the observables are themselves persistent over generations, how do they influence the peristence in economic outcomes in turn. Below we address this question.

Denoting the data matrix of the log of individual earnings, other income and consumption as y_{ft}, we proceed as follows:

1. We regress the log of each outcome variable, y_{ft}, on a full set of year and cohort dummies, and denote estimated residuals as $\hat{y}_{ft}^{(1)}$. These are our baseline outcome measures.

2. Next, we regress our baseline outcomes $\hat{y}_{ft}^{(1)}$ on a set of observables x_{ft}. That is, we estimate least square projections:[94]

$$\hat{y}_{ft}^{(1)} = \beta x_{ft} + \varepsilon_{ft}. \tag{C.5}$$

3. From the previous step we recover predicted values, as well as residuals. Specifically, we define:

$$\hat{y}_{ft}^{(2)} \equiv \hat{\beta} x_{ft} \tag{C.6}$$

and

$$\hat{y}_{ft}^{(3)} \equiv \hat{\varepsilon}_{ft} \tag{C.7}$$

[94] The matrix of controls x_{ft} includes dummies for family size, number of children, state of residence, employment status, race and education.

For each of the measures $\hat{y}_{ft}^{(i)}$ we compute a set of variances and covariances. Each set of second moments can then be used to separately estimate structural model parameters.

4. We estimate the GMM model separately for each set of variance-covariance moments of $\hat{y}_{ft}^{(i)}$ $(i \in 1, 2, 3)$. This delivers different sets of parameter estimates. Comparing these estimates is helpful to establish whether the transmission of inequality is due to observable or unobservable components.

Table C.3: Variances for Parents and Children

Variable	**Generation**	**Baseline** (1)	**Observable** (2)	**Unobservable** (3)
Earnings	Parent	0.291	0.093	0.182
	Child	0.248	0.057	0.177
Other Income	Parent	0.808	0.084	0.696
	Child	0.534	0.081	0.441
Consumption	Parent	0.097	0.024	0.066
	Child	0.114	0.024	0.087

Table C.4: Baseline Estimates: Intergenerational Elasticity

Variables	**Parameters**	**Baseline** (1)	**Observable** (2)	**Unobservable** (3)
Earnings	γ	0.230	0.339	0.109
		(0.027)	(0.022)	(0.027)
Other Income	ρ	0.100	0.248	0.021
		(0.023)	(0.039)	(0.029)
$\bar{e}_f^p$ on $n_{f,t}^k$	λ	0.206	0.255	0.060
		(0.032)	(0.028)	(0.038)
$\bar{n}_f^p$ on $e_{f,t}^k$	θ	0.055	0.111	0.003
		(0.019)	(0.028)	(0.017)
Consumption Shifters	ϕ	0.154	0.450	0.010
		(0.032)	(0.042)	(0.034)
No. of Parent-Child Pairs	N	760	760	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. *Baseline* refers to data that is purged of year and birth cohort effects (viz., $\hat{y}_{ft}^{(1)}$ in equation C.5). These data are then regressed on various controls (namely, dummies for family size, state of residence, number of children, employment status, race and education). *Observable* refers to the fitted values from this regression (viz., $\hat{y}_{ft}^{(2)}$ in equation C.6), while *Unobservable* refers to its residual (viz., $\hat{y}_{ft}^{(3)}$ in equation C.7). The average age for parents in the sample is 47 years; that of children is 37 years.

Table C.3 reports the cross-sectional variances of earnings, other income and consumption for parents and for children. Columns 1-3 correspond to equations (C.5), (C.6) and (C.7): column 1 reports the variance controlling only for time and cohort effects, as in equation equation (C.5); column 2 reports the fitted variance, defined as the variance explained by observables, as in equation (C.6); and column 3 reports the variance of the residual, equation (C.7).

Table C.5: Baseline Estimates: Variances and Covariances of Idiosyncratic Components

	Parameters	**Baseline** (1)	**Observable** (2)	**Unobservable** (3)
Parental Outcomes: Variances				
Permanent Earnings	$\sigma^2_{\bar{e}^p}$	0.295	0.095	0.182
		(0.021)	(0.006)	(0.011)
Permanent Other Income	$\sigma^2_{\bar{n}^p}$	0.806	0.084	0.696
		(0.06)	(0.01)	(0.059)
Permanent Consumption Shifters	$\sigma^2_{\bar{q}^p}$	1.031	0.196	0.789
		(0.065)	(0.022)	(0.064)
Child Idiosyncratic Heterogeneity: Variances				
Permanent Earnings	$\sigma^2_{\delta^k}$	0.228	0.041	0.175
		(0.011)	(0.002)	(0.01)
Permanent Other Income	$\sigma^2_{\varepsilon^k}$	0.511	0.062	0.440
		(0.043)	(0.004)	(0.031)
Permanent Consumption Shifters	$\sigma^2_{\psi^k}$	0.730	0.104	0.584
		(0.056)	(0.008)	(0.039)
Parental Outcomes: Covariances				
Consumption Shifters & Earnings	$\sigma_{\bar{q}^p,\bar{e}^p}$	-0.271	-0.120	-0.120
		(0.024)	(0.01)	(0.016)
Consumption Shifters & Other Income	$\sigma_{\bar{q}^p,\bar{n}^p}$	-0.818	-0.116	-0.669
		(0.061)	(0.015)	(0.062)
Earnings and Other Income	$\sigma_{\bar{e}^p,\bar{n}^p}$	0.070	0.059	-0.012
		(0.013)	(0.007)	(0.012)
Child Idiosyncratic Heterogeneity: Covariances				
Consumption Shifters & Earnings	σ_{ψ^k,δ^k}	-0.247	-0.058	-0.165
		(0.018)	(0.004)	(0.016)
Consumption Shifters & Other Income	$\sigma_{\psi^k,\varepsilon^k}$	-0.522	-0.068	-0.430
		(0.048)	(0.005)	(0.034)
Earnings & Other Income	$\sigma_{\delta^k,\varepsilon^k}$	0.075	0.031	0.036
		(0.013)	(0.002)	(0.013)
No. of Parent-Child Pairs	N	760	760	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. *Baseline* refers to data that is purged of year and birth cohort effects (viz., $\hat{y}^{(1)}_{ft}$ in equation C.5). These data are then regressed on various controls (namely, dummies for family size, state of residence, number of children, employment status, race, and education). *Observable* refers to the fitted value from this regression (viz., $\hat{y}^{(2)}_{ft}$ in equation C.6), while *Unobservable* refers to its residual (viz., $\hat{y}^{(3)}_{ft}$ in equation C.7).

Next, we use these variances and other covariances amongst the economic outcomes to estimate the parameters for intergenerational eslaticity (reported in Table C.4) and for the variance-covariance structure of the idiosyncratic shocks specific to a particular generation (reported in Table C.5). From Table C.4 is is clear that all pass-through parameters in the baseline estimation are primarily driven by persistence in observables, while only earnings has some part that is explained by unobservable factors that are linked across generations.

Table C.6: Share of Child Inequality Explained by Parental Heterogeneity

Variables	**Baseline** (1)	**Observable** (2)	**Unobservable** (3)
Earnings	8.0	28.7	1.2
Other Income	4.4	23.4	0.2
Consumption	29.9	32.8	5.6

Note: Values represent the percentage share of cross-sectional variances for younger generation that is explained by parental factors. Numbers obtained using parameter estimates from Tables C.4 and C.5.

Role of Education

Table C.7: Baseline Estimates: Intergenerational Elasticity for Observables

	Parameters	**Observable** (2)	**Education** (3)	**Other** (4)
Earnings	γ	0.339	0.258	0.304
		(0.022)	(0.035)	(0.023)
Other Income	ρ	0.248	0.188	0.207
		(0.039)	(0.027)	(0.05)
$\bar{e}^p_f$ on $n^k_{f,t}$	λ	0.255	0.183	0.271
		(0.028)	(0.018)	(0.041)
$\bar{n}^p_f$ on $e^k_{f,t}$	θ	0.111	0.189	0.054
		(0.028)	(0.033)	(0.028)
Consumption Shifters	ϕ	0.450	0.410	0.354
		(0.042)	(0.029)	(0.062)
No. of Parent-Child Pairs	N	760	760	760

Note: Bootstrap standard errors with 100 repetitions are reported in parentheses. *Observable* refers to the total fitted value of the regression of the data (purged off of year and birth cohort effects) on dummies for family size, state of residence, number of children, employment status, race and education. *Education* refers to the fitted value of the regression of the data on education only, while *Other* refers to the fitted value of the other observable control variables. The average age for parents is 47 years, while that for children is 37 years in the sample.

Table C.8: Mobility Matrix for Education

Child \ Parent	<12 years	High School	College Dropout	College & above
<12 years	**21.83**	4.93	0	0
High School	40.47	**39.95**	19.27	7.80
College Dropout	20.93	25.58	**42.52**	14.99
College & above	16.77	29.54	38.21	**77.20**

C.5.2 The Impact of Parental Factors on Inequality

Variance Accounting Calculations. As reported in Section 4.4.1, the relative contribution of parental factors in the cross-sectional variance of earnings among their kids' generation can be computed as the ratio

$$\frac{Var\left[e^k(p)\right]}{Var\left[e^k\right]} = \frac{\gamma^2\sigma^2_{\bar{e}^p} + \theta^2\sigma^2_{\bar{n}^p} + 2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p}}{\sigma^2_{\delta^k} + \gamma^2\sigma^2_{\bar{e}^p} + \theta^2\sigma^2_{\bar{n}^p} + 2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p}}. \tag{C.8}$$

Then, substituting the parameter estimates from Tables C.4 and C.5 in equation C.8, one can obtain the estimates in the first row of Table C.6. That is, we can write:

$$\begin{aligned} & \frac{\gamma^2\sigma^2_{\bar{e}^p} + \theta^2\sigma^2_{\bar{n}^p} + 2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p}}{\sigma^2_{\delta^k} + \gamma^2\sigma^2_{\bar{e}^p} + \theta^2\sigma^2_{\bar{n}^p} + 2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p}} \\ = \; & \frac{(0.230^2)(0.295) + (0.055^2)(0.806) + 2(0.230)(0.055)(0.070)}{0.228 + (0.230^2)(0.295) + (0.055^2)(0.806) + 2(0.230)(0.055)(0.070)} = 8.0\%. \end{aligned}$$

Similarly, the contribution of parental factors to the cross-sectional variances of other income and consumption in the children's generation is given by the ratios,

$$\frac{Var\left[n^k(p)\right]}{Var\left[n^k\right]} \tag{C.9}$$

and

$$\frac{Var\left[c^k(p)\right]}{Var\left[c^k\right]} \tag{C.10}$$

where

$$Var\left[n^k(p)\right] = \rho^2\sigma^2_{\bar{n}^p} + \lambda^2\sigma^2_{\bar{e}^p} + 2\rho\lambda\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.11}$$

$$Var\left[n^k\right] = Var\left[n^k(p)\right] + \sigma^2_{\varepsilon^k} \tag{C.12}$$

$$\begin{aligned} Var\left[c^k(p)\right] = \; & \phi^2\sigma^2_{\bar{q}^p} + (\gamma+\lambda)^2\sigma^2_{\bar{e}^p} + (\rho+\theta)^2\sigma^2_{\bar{n}^p} \\ & + 2\left[(\gamma+\lambda)\phi\sigma_{\bar{e}^p,\bar{q}^p} + (\rho+\theta)\phi\sigma_{\bar{n}^p,\bar{q}^p} + (\rho+\theta)(\gamma+\lambda)\sigma_{\bar{e}^p,\bar{n}^p}\right] \end{aligned} \tag{C.13}$$

$$Var\left[c^k\right] = Var\left[c^k(p)\right] + \sigma^2_{\varepsilon^k} + \sigma^2_{\psi^k} + \sigma^2_{\delta^k} + 2\left(\sigma_{\psi^k,\varepsilon^k} + \sigma_{\psi^k,\delta^k} + \sigma_{\delta^k,\varepsilon^k}\right). \tag{C.14}$$

Counterfactual Distributions. In order to compare the actual distribution of outcomes for children with the counterfactual distributions where parental effects are shut down, we assume that the permanent parental and idiosyncratic child components of earnings, other income and consumption jointly follow a Gaussian distribution in logarithms[95]:

$$\begin{pmatrix} \bar{e}^p_f \\ \bar{n}^p_f \\ \bar{q}^p_f \\ \delta^k_f \\ \varepsilon^k_f \\ \psi^k_f \end{pmatrix} \sim \mathbf{N}\left[\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2_{\bar{e}^p} & \sigma_{\bar{e}^p,\bar{n}^p} & \sigma_{\bar{e}^p,\bar{q}^p} & 0 & 0 & 0 \\ \sigma_{\bar{e}^p,\bar{n}^p} & \sigma^2_{\bar{n}^p} & \sigma_{\bar{n}^p,\bar{q}^p} & 0 & 0 & 0 \\ \sigma_{\bar{e}^p,\bar{q}^p} & \sigma_{\bar{n}^p,\bar{q}^p} & \sigma^2_{\bar{q}^p} & 0 & 0 & 0 \\ 0 & 0 & 0 & \sigma^2_{\delta^k} & \sigma_{\delta^k,\varepsilon^k} & \sigma_{\delta^k,\psi^k} \\ 0 & 0 & 0 & \sigma_{\delta^k,\varepsilon^k} & \sigma^2_{\varepsilon^k} & \sigma_{\psi^k,\varepsilon^k} \\ 0 & 0 & 0 & \sigma_{\psi^k,\delta^k} & \sigma_{\psi^k,\varepsilon^k} & \sigma^2_{\psi^k} \end{pmatrix}\right]$$

Then, by the property of a joint Normal distribution, any linear combination of the constituent random variables also follows a Normal distribution. For example, we can assume that the idiosyncratic part of permanent child consumption, $\left(\varepsilon^k_f + \psi^k_f + \delta^k_f\right)$, follows a Normal distribution with zero mean and variance equal to $\sigma^2_{\varepsilon^k} + \sigma^2_{\psi^k} + \sigma^2_{\delta^k} + 2\left(\sigma_{\psi^k,\varepsilon^k} + \sigma_{\psi^k,\delta^k} + \sigma_{\delta^k,\varepsilon^k}\right)$. Such child idiosyncratic components are by definition independent of any parental influence, and hence can be used to generate the counterfactual distribution for the children. Now, since the logarithmic random variables follow the Gaussian distribution (by assumption), they will follow the Lognormal distribution in their levels. Figure 4.2 reports the difference in the probability density functions with and without parental influence.

C.6 Evolution of Inequality

What degree of persistence would generate, all else equal, growing dispersion across generations? To answer this question, one needs to derive a threshold value of persistence as a function of the inequality in that generation. In order to get a closed form expression for these threshold values of persistence, we shut down the cross-persistence terms, that is, restrict $\lambda = \theta = 0$. With these parameter restrictions, earnings evolve through generations of the same family according to:

$$\begin{aligned} e^{k_1} &= \gamma\bar{e}^p + \delta^{k_1} \\ e^{k_2} &= \gamma^2\bar{e}^p + \gamma\delta^{k_1} + \delta^{k_2} \\ &\vdots \\ e^{k_t} &= \gamma^t\bar{e}^p + \sum_{j=1}^{t}\gamma^{t-j}\delta^{k_t} \end{aligned}$$

where the superscript $\{k_t\}$ identifies the t^{th} generation of kids. Since $\gamma \in (0,1)$, there exists a long run stationary distribution for earnings. Assuming $\text{Var}(\delta^{k_t}) = \sigma^2_{\delta^k}\ \forall t$ and $\text{Cov}(\delta^{k_t}, \delta^{k_{t'}}) = 0\ \forall t \neq t'$,

[95] The mean of the logarithmic variables are zero because we consider de-meaned variables net of year and cohort fixed effects.

the variance of the stationary distribution of e, denoted by $\text{Var}(e^*)$, is

$$\text{Var}(e^*) = \lim_{t\to\infty}\left[\gamma^{2t}\sigma^2_{\bar{e}^p} + \sum_{j=1}^{t}\gamma^{2(t-j)}\sigma^2_{\delta^k}\right] = \frac{\sigma^2_{\delta^k}}{1-\gamma^2} \tag{C.15}$$

Similarly, one can derive the stationary variances for other income and consumption as,

$$\text{Var}(n^*) = \frac{\sigma^2_{\varepsilon^k}}{1-\rho^2} \tag{C.16}$$

$$\text{Var}(c^*) = \frac{\sigma^2_{\psi^k}}{1-\phi^2} + \frac{\sigma^2_{\delta^k}}{1-\gamma^2} + \frac{\sigma^2_{\varepsilon^k}}{1-\rho^2} + \frac{2\sigma_{\delta^k,\varepsilon^k}}{1-\gamma\rho} + \frac{2\sigma_{\psi^k,\varepsilon^k}}{1-\phi\rho} + \frac{2\sigma_{\psi^k,\delta^k}}{1-\phi\gamma}. \tag{C.17}$$

Table C.9: Intergenerational Elasticities

	Parameters	**Estimates** (1)
Earnings	γ	0.279
		(0.048)
Other Income	ρ	0.020
		(0.041)
Consumption Shifters	ϕ	0.006
		(0.047)
No. of Parent-Child Pairs	N	403

Note: Bootstrap standard errors (100 repetitions) in parentheses. Parental and child ages vary between 30 and 40. Parameters λ and θ are set to zero. Average parental age is 37 years, while average age of children is 35. Food expenditures are used as a measure of consumption. Estimates use cross-sectional data variation net of cohort and year effects.

Plugging in estimated values for the parameters in equations (C.15) through (C.17),[96] one can identify the threshold values of the persistence parameters beyond which there will be rising inequality. Using equation (C.15), we identify the threshold value of γ above which the variance of earnings would grow from the value estimated in the parents' generation: this is the value of γ such that $\text{Var}(e^*) \geq \text{Var}(e^p)$. This threshold value of γ is given by $\gamma^p \equiv \sqrt{1 - \frac{\sigma^2_{\delta^k}}{\text{Var}(e^p)}}$. Any γ larger than γ^p implies growing earnings variance. Based on the parameter estimates in Tables C.9 and C.10, $\sigma^2_{\delta^k} = 0.246 > \text{Var}(e^p) = 0.183$, making γ^p an imaginary number. This essentially implies that any non-negative value of γ would result in increasing earnings inequality from the level in the parents' generation. Since our estimate of the current value of γ $(= 0.279)$ is positive, the model implies that the earnings variance should become larger in the next generation k_1. In fact, earnings variance in the child generation, $\text{Var}\left(e^{k_1}\right) = 0.261$ is larger than in the parents' one,

[96]Since we restrict the parameters $\lambda = \theta = 0$, we need to re-estimate our baseline model with this additional restriction. Additionally, we restrict the age range between 30 and 40 years for both parents and kids, in order to facilitate comparison of inequality across different generations in the same age range. These estimates are reported in Tables C.9 and C.10.

$\text{Var}\left(e^p\right) = 0.183.$

Table C.10: Idiosyncratic Variances & Covariances

	Parameters	Estimates (1)
Parental Outcomes: Variances		.
Permanent Earnings	$\sigma^2_{\bar{e}^p}$	0.183 (0.012)
Permanent Other Income	$\sigma^2_{\bar{n}^p}$	0.877 (0.128)
Permanent Consumption Shifters	$\sigma^2_{\bar{q}^p}$	0.956 (0.134)
Child Idiosyncratic Heterogeneity: Variances		.
Permanent Earnings	$\sigma^2_{\delta^k}$	0.246 (0.013)
Permanent Other Income	$\sigma^2_{\varepsilon^k}$	0.630 (0.038)
Permanent Consumption Shifters	$\sigma^2_{\psi^k}$	0.848 (0.037)
Parental Outcomes: Covariances		.
Consumption Shifters & Earnings	$\sigma_{\bar{q}^p,\bar{e}^p}$	-0.122 (0.029)
Consumption Shifters & Other Income	$\sigma_{\bar{q}^p,\bar{n}^p}$	-0.841 (0.126)
Earnings and Other Income	$\sigma_{\bar{e}^p,\bar{n}^p}$	-0.000 (0.025)
Child Idiosyncratic Heterogeneity: Covariances		.
Consumption Shifters & Earnings	σ_{ψ^k,δ^k}	-0.247 (0.020)
Consumption Shifters & Other Income	$\sigma_{\psi^k,\varepsilon^k}$	-0.620 (0.032)
Earnings & Other Income	$\sigma_{\delta^k,\varepsilon^k}$	0.056 (0.017)
No. of Parent-Child Pairs	N	403

Note: Bootstrap standard errors (100 repetitions) in parentheses. This table uses the same sample and model specification as Table C.9.

Starting from the children generation, and using equation (C.15) again, we can find the threshold value of γ above which the earnings variance after the child generation would be growing; that is,

$$\gamma^{k_1} \equiv \sqrt{1 - \frac{\sigma^2_{\delta^k}}{\text{Var}(e^{k_1})}} = \sqrt{1 - \frac{0.246}{0.261}} = 0.24.$$

This is plotted as the dashed vertical line in Figure C.3. Any value of γ to the right of that vertical line implies growing earnings variance. Since our estimate of γ (= 0.279) lies to the

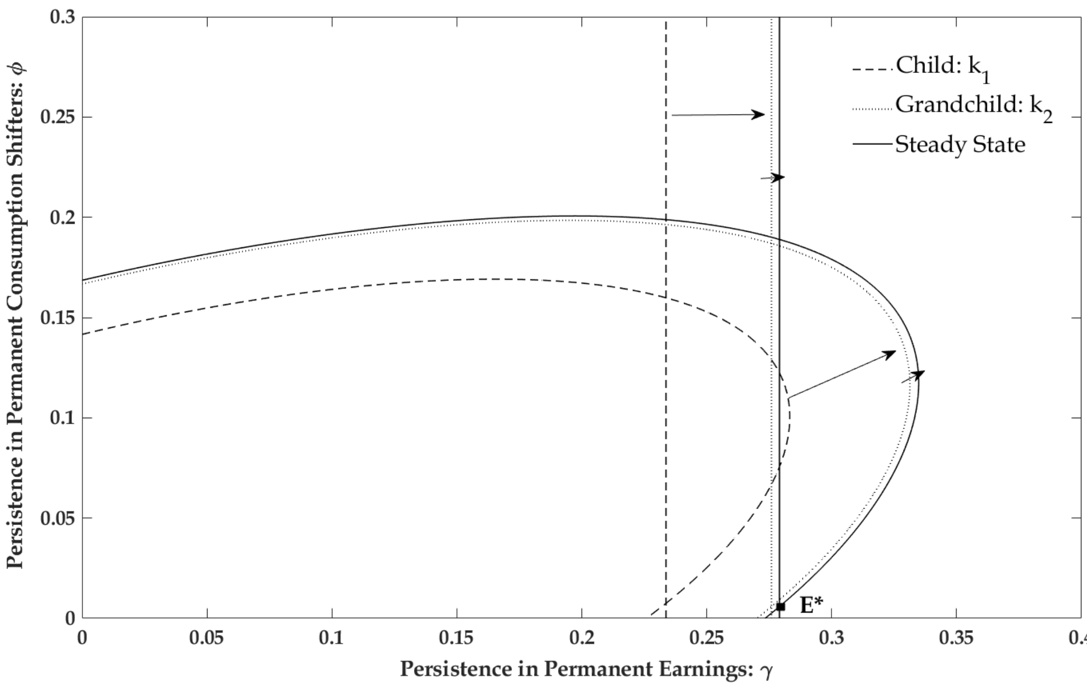

Figure C.3: Implication of γ and ϕ for Long Run Inequality (Age: 30-40)

right of the new threshold γ^{k_1}, the threshold corresponding to the generation of grandchildren k_2 (denoted by the dotted vertical line in Figure C.3) will lie further to the right of γ^{k_1}; one can repeat these calculations over and over again.[97] Eventually, the economy settles down at the stationary distribution of earnings where the threshold is defined as

$$\gamma^* \equiv \sqrt{1 - \frac{\sigma^2_{\delta^k}}{\mathrm{Var}(e^*)}} = 0.279,$$

which is the estimated level of γ.

We can perform a similar exercise for the evolution of the variance of consumption using equation (C.17). Instead of a single persistence parameter γ, as in the case of earnings, the variance of consumption is a function of three persistence parameters: γ, ρ and ϕ. To make interpretation easier, we hold ρ constant at its estimated value and study the thresholds of γ and ϕ that imply increasing or decreasing consumption variance. Equation (C.17) shows that $\mathrm{Var}(c^*)$ is a non-linear function of γ and ϕ. First we ask what combinations of γ and ϕ imply that the variance of consumption is increasing across subsequent generations. For that we would like to plot the threshold value,

$$\mathrm{Var}(c^g) = \frac{\sigma^2_{\psi^k}}{1-\phi^2} + \frac{\sigma^2_{\delta^k}}{1-\gamma^2} + \frac{\sigma^2_{\varepsilon^k}}{1-\rho^2} + \frac{2\sigma_{\delta^k,\varepsilon^k}}{1-\gamma\rho} + \frac{2\sigma_{\psi^k,\varepsilon^k}}{1-\phi\rho} + \frac{2\sigma_{\psi^k,\delta^k}}{1-\phi\gamma},$$

[97] We find $\gamma^{k_2} = 0.276$, which is larger than γ^{k_1} but still slightly smaller than 0.279.

for each generation $g = \{p, k_1, k_2, ...\}$ as a function of γ and ϕ, holding all other parameters constant. However, there is no combination of γ and ϕ in the economically meaningful range $[0, 1]$ that satisfies the threshold value equation for $\text{Var}(c^p)$. Therefore, any point in the $(\gamma, \phi) \in [0, 1]^2$ space will imply rising consumption inequality from the parents' generation. This finding is corroborated by the fact that $\text{Var}\left(c^{k_1}\right) = 0.117 > \text{Var}(c^p) = 0.09$.

Next, we plot the threshold starting from the children's generation, denoted by the dashed ellipse in Figure C.3. Since the estimated point, labelled E^*, with values $(\gamma, \phi) = (0.28, 0.01)$, lies outside this ellipse, the grandchildren's generation should have a larger consumption variance than the children's generation. Indeed, plotting the corresponding threshold for the grandchild generation (denoted by the dotted ellipse in Figure C.3), we find that it lies outside that for the children with $\text{Var}(c^{k_2}) = 0.124 > \text{Var}(c^{k_1}) = 0.117$. These dynamics are replicated across generations until the economy settles at the stationary distribution of consumption which gives rise to the solid elliptical threshold of γ and ϕ in Figure C.3.[98]

While the analysis above shows how the estimates of current parameter values help make sense of the evolution of earnings and consumption variances across generations, these hypothetical dynamics are specific to the parameter estimates we feed into the model, which are in turn determined by the raw data moments that we currently observe. For example, the dynamics of increasing earnings variance are contingent on whether our raw data imply $\text{Var}(e^p) < \text{Var}(e^k)$. As an example of an alternative scenario, we use the estimates in column (2) of Tables 4.11 and 4.12 which does not restrict the age to be between 30 and 40 years, but keeps the $\lambda = \theta = 0$ restriction. Relaxing our age restriction implies $\text{Var}(e^p) > \text{Var}(e^k)$, so that the thresholds of γ approach the long run threshold from the right, rather than from the left as in Figure C.3, suggesting decreasing earnings variance across generations. Similarly, the dynamics of consumption and other income inequality in the long run are also dictated by the empirically observed moments.

Relaxing Age Restriction. We replicate the above analysis of inequality evolution using a parametrization of the model based on a sample without age restrictions. This means that the relevant parameter estimates are obtained from columns (2) of Tables 4.11 and 4.12.

The threshold value of γ beyond which the earnings inequality is increasing in the parents' generation is given by

$$\gamma^p \equiv \sqrt{1 - \frac{\sigma^2_{\delta^k}}{\text{Var}(e^p)}} = 0.506,$$

and is shown as the dot-dashed vertical line in Figure C.4. Since the estimate of the current value of γ $(= 0.340)$ lies to the left of that line, the model implies that the earnings variance should become smaller in the next generation k_1. We corroborate this using equation (C.15) again to find the threshold value of γ above which the earnings variance in the child generation should be growing.

[98]The stationary locus for earnings (the solid vertical line) and that of consumption (the solid ellipse) intersect at two points. One of those points, denoted by E^*, corresponds to the GMM point estimate of γ and ϕ. The other intersection point cannot be an equilibrium of the model because the stationary locus for other income (not plotted here) passes only through E^*.

We find

$$\gamma^{k_1} \equiv \sqrt{1 - \frac{\sigma^2_{\delta^k}}{\mathrm{Var}(e^{k1})}} = 0.367,$$

which is less than γ^p. Once again the estimated value of $\gamma = 0.340$ lies to the left of this new threshold γ^{k1}, and so the threshold corresponding to the generation of grandchildren k_2 will lie further to the left of γ^{k_1}, and so on. Eventually, the economy settles down at the stationary distribution of earnings where the threshold is defined as $\gamma^* \equiv \sqrt{1 - \frac{\sigma^2_{\delta^k}}{\mathrm{Var}(e^*)}} = 0.340$, which is the estimated level of γ.

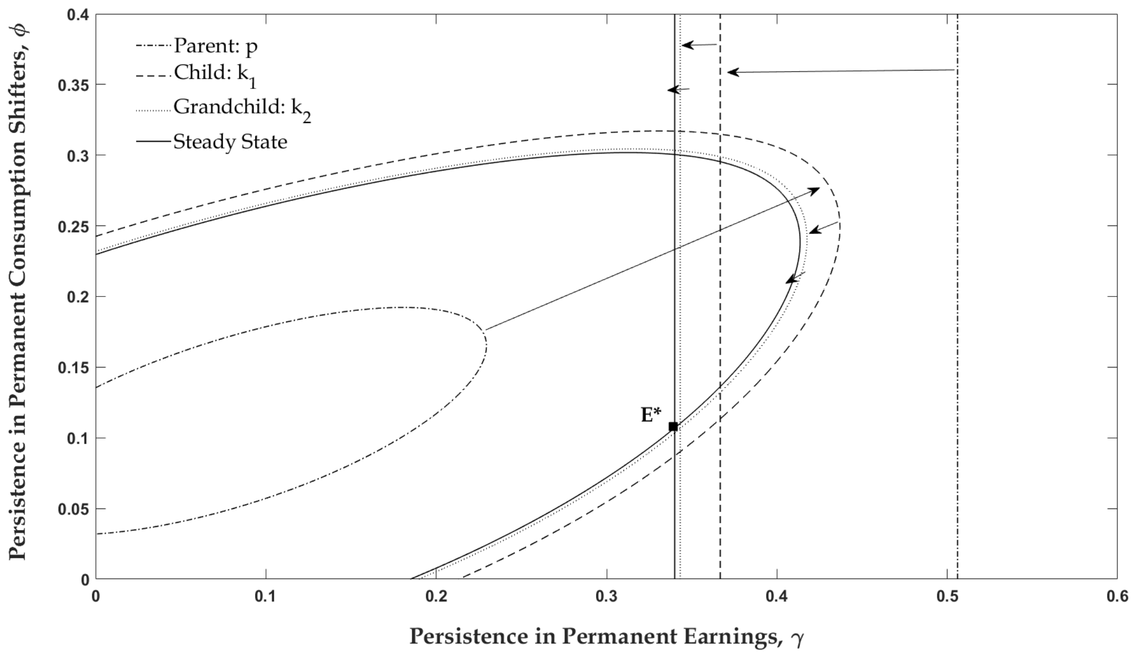

Figure C.4: Implication of γ and ϕ for Long Run Earnings & Consumption Inequality

We again perform a similar exercise for the consumption variance using equation (C.17). The variance of consumption is a function of three persistence parameters: γ, ρ and ϕ. We hold ρ constant at its estimated value and study the thresholds of γ and ϕ that imply increasing or decreasing consumption variance. First we ask what combinations of γ and ϕ imply that the variance of consumption is increasing across generations. For that we plot the threshold value

$$\mathrm{Var}(c^p) = \frac{\sigma^2_{\psi^k}}{1-\phi^2} + \frac{\sigma^2_{\delta^k}}{1-\gamma^2} + \frac{\sigma^2_{\varepsilon^k}}{1-\rho^2} + \frac{2\sigma_{\delta^k,\varepsilon^k}}{1-\gamma\rho} + \frac{2\sigma_{\psi^k,\varepsilon^k}}{1-\phi\rho} + \frac{2\sigma_{\psi^k,\delta^k}}{1-\phi\gamma},$$

as a function of γ and ϕ. This is shown as the dot-dashed ellipse in Figure C.4. Any point inside that ellipse implies the variance of consumption for the child generation is less than their parents. Since the estimated point, labelled E^*, with values $(\gamma, \phi) = (0.340, 0.107)$, lies outside this ellipse, the children's generation should have a larger consumption variance than the parental generation. Indeed, plotting the corresponding threshold for the child generation, (denoted by the outermost dashed ellipse in Figure C.4), we find that it lies outside that for the parents with $\mathrm{Var}(c^{k_1}) = 0.114 > \mathrm{Var}(c^p) = 0.096$. However, our estimate values of $(\gamma, \phi) = (0.340, 0.107)$ lie inside the ellipse for the child generation. This means that the generation of grandchildren k_2 should exhibit lower consumption variance than the child generation k_1, and therefore should have a threshold ellipse which lies inside that for the child generation. These dynamics are replicated across generations until the economy settles at the stationary distribution of consumption which gives rise to the solid black elliptical threshold of γ and ϕ in Figure C.4.

C.7 Robustness and Extensions

In this appendix we present additional empirical results for the different extensions and robustness checks of the baseline specification that we considered in Section 4.6.

C.7.1 Estimates by Child Birth-Cohort

Table C.11: Parental Importance by Child-Cohort (Age: 30-40)

Variables	**All Cohorts** (1)	**1960s Cohort** (2)	**1970s Cohort** (3)
Earnings	4.0	4.4	5.3
Other Income	1.6	1.3	2.9
Consumption	24.4	37.7	15.9

Note: All numbers are percentages (%) and are based on parameter estimates in Tables 4.10 and C.12.

Table C.12: Estimates by Child Cohort: Idiosyncratic Components (Age: 30-40)

	Parameters	**All Cohorts** (1)	**1960s Cohort** (2)	**1970s Cohort** (3)
Parental Outcomes: Variances				
Permanent Earnings	$\sigma^2_{\bar{e}^p}$	0.199	0.172	0.225
		(0.019)	(0.021)	(0.026)
Permanent Other Income	$\sigma^2_{\bar{n}^p}$	0.845	0.945	0.752
		(0.105)	(0.16)	(0.157)
Permanent Consumption Shifters	$\sigma^2_{\bar{q}^p}$	0.911	0.977	0.840
		(0.115)	(0.157)	(0.153)
Child Idiosyncratic Heterogeneity: Variances				
Permanent Earnings	$\sigma^2_{\delta^k}$	0.241	0.232	0.245
		(0.017)	(0.021)	(0.025)
Permanent Other Income	$\sigma^2_{\varepsilon^k}$	0.658	0.561	0.747
		(0.067)	(0.1)	(0.162)
Permanent Consumption Shifters	$\sigma^2_{\psi^k}$	0.869	0.816	0.900
		(0.075)	(0.129)	(0.196)
Parental Outcomes: Covariances				
Consumption Shifters & Earnings	$\sigma_{\bar{q}^p,\bar{e}^p}$	-0.126	-0.060	-0.187
		(0.037)	(0.031)	(0.041)
Consumption Shifters & Other Income	$\sigma_{\bar{q}^p,\bar{n}^p}$	-0.798	-0.887	-0.711
		(0.106)	(0.152)	(0.149)
Earnings and Other Income	$\sigma_{\bar{e}^p,\bar{n}^p}$	-0.006	-0.044	0.029
		(0.029)	(0.029)	(0.029)
Child Idiosyncratic Heterogeneity: Covariances				
Consumption Shifters & Earnings	σ_{ψ^k,δ^k}	-0.232	-0.269	-0.189
		(0.028)	(0.036)	(0.039)
Consumption Shifters & Other Income	$\sigma_{\psi^k,\varepsilon^k}$	-0.654	-0.583	-0.714
		(0.069)	(0.114)	(0.181)
Earnings & Other Income	$\sigma_{\delta^k,\varepsilon^k}$	0.047	0.078	0.013
		(0.025)	(0.026)	(0.044)
No. of Parent-Child Pairs	N	336	166	170

Note: Bootstrap standard errors with 100 repetitions are reported in parentheses. This table uses the same sample and model specification as Table 4.10.

C.7.2 Estimates under Alternative Definitions of 'Other Income'

Table C.13: Decomposition of Other Income: Idiosyncratic Components

	Parameters	**Just Transfers** (1)	**Spouse Earnings** (2)	**Other Income** (3)
Parental Outcomes: Variances				
Permanent Earnings	$\sigma^2_{\bar{e}^p}$	0.287	0.295	0.295
		(0.027)	(0.027)	(0.027)
Permanent Other Income	$\sigma^2_{\bar{n}^p}$	1.297	0.294	0.459
		(0.128)	(0.021)	(0.041)
Permanent Consumption Shifters	$\sigma^2_{\bar{q}^p}$	1.504	0.502	0.650
		(0.132)	(0.042)	(0.055)
Child Idiosyncratic Heterogeneity: Variances				
Permanent Earnings	$\sigma^2_{\delta^k}$	0.213	0.196	0.205
		(0.015)	(0.011)	(0.012)
Permanent Other Income	$\sigma^2_{\varepsilon^k}$	1.063	0.296	0.441
		(0.085)	(0.018)	(0.038)
Permanent Consumption Shifters	$\sigma^2_{\psi^k}$	1.292	0.460	0.589
		(0.096)	(0.029)	(0.042)
Parental Outcomes: Covariances				
Consumption Shifters & Earnings	$\sigma_{\bar{q}^p,\bar{e}^p}$	-0.225	-0.259	-0.247
		(0.042)	(0.029)	(0.026)
Consumption Shifters & Other Income	$\sigma_{\bar{q}^p,\bar{n}^p}$	-1.314	-0.302	-0.458
		(0.125)	(0.027)	(0.045)
Earnings and Other Income	$\sigma_{\bar{e}^p,\bar{n}^p}$	0.044	0.063	0.051
		(0.035)	(0.015)	(0.017)
Child Idiosyncratic Heterogeneity: Covariances				
Consumption Shifters & Earnings	σ_{ψ^k,δ^k}	-0.222	-0.199	-0.204
		(0.033)	(0.017)	(0.021)
Consumption Shifters & Other Income	$\sigma_{\psi^k,\varepsilon^k}$	-1.081	-0.289	-0.421
		(0.103)	(0.02)	(0.035)
Earnings & Other Income	$\sigma_{\delta^k,\varepsilon^k}$	0.059	0.056	0.050
		(0.028)	(0.014)	(0.017)
No. of Parent-Child Pairs	N	459	459	459

Note: Bootstrap standard errors with 100 repetitions are reported in parentheses. This table uses the same sample and model specification as Table 4.5.

Table C.14: Estimated Variances of Components of *Other Income*

Variable	Generation	Just Transfers (1)	Spouse Earnings (2)	Other Income (3)
Earnings	Parent	0.287	0.295	0.295
	Child	0.229	0.229	0.229
***'Other Income'* Component**	Parent	1.297	0.294	0.459
	Child	1.068	0.322	0.457
Consumption	Parent	0.098	0.094	0.096
	Child	0.113	0.113	0.113

Note: This table uses parameter estimates from Tables 4.5 and C.13.

C.7.3 Model using Panel Data

In this appendix we present the full set of moment conditions for the model using panel data variation, and the identification strategy for all the parameters. We also report the estimates of the variances of the transitory shocks, that were averaged out in the baseline specification with only cross-sectional variation.

Parent Variance

$$Var\left(e^p_{f,t}\right) = \sigma^2_{\bar{e}^p} + \sigma^2_{\zeta^p} \tag{C.18}$$

$$Var\left(n^p_{f,t}\right) = \sigma^2_{\bar{n}^p} + \sigma^2_{u^p} \tag{C.19}$$

$$\begin{aligned} Var\left(c^p_{f,t}\right) &= \sigma^2_{\bar{q}^p} + \sigma^2_{\bar{e}^p} + \sigma^2_{\bar{n}^p} + \sigma^2_{v^p} \\ &+ 2\left(\sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) + \left[f\left(r\right)\right]^2\left(\sigma^2_{u^p} + \sigma^2_{\zeta^p} + 2\sigma_{\zeta^p,u^p}\right) \end{aligned} \tag{C.20}$$

Child Variance

$$Var\left(e^k_{f,t}\right) = \gamma^2\sigma^2_{\bar{e}^p} + \theta^2\sigma^2_{\bar{n}^p} + \sigma^2_{\delta^k} + \sigma^2_{\zeta^k} + 2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.21}$$

$$Var\left(n^k_{f,t}\right) = \rho^2\sigma^2_{\bar{n}^p} + \lambda^2\sigma^2_{\bar{e}^p} + \sigma^2_{\varepsilon^k} + \sigma^2_{u^k} + 2\rho\lambda\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.22}$$

$$\begin{aligned} Var\left(c^k_{f,t}\right) &= \phi^2\sigma^2_{\bar{q}^p} + \left(\gamma+\lambda\right)^2\sigma^2_{\bar{e}^p} + \left(\rho+\theta\right)^2\sigma^2_{\bar{n}^p} + \sigma^2_{\varepsilon^k} + \sigma^2_{\psi^k} + \sigma^2_{\delta^k} \\ &+ 2\left[\left(\gamma+\lambda\right)\phi\sigma_{\bar{q}^p,\bar{e}^p} + \left(\rho+\theta\right)\phi\sigma_{\bar{q}^p,\bar{n}^p} + \left(\rho+\theta\right)\left(\gamma+\lambda\right)\sigma_{\bar{e}^p,\bar{n}^p}\right] \\ &+ 2\left[\sigma_{\psi^k,\varepsilon^k} + \sigma_{\psi^k,\delta^k} + \sigma_{\delta^k,\varepsilon^k}\right] \\ &+ \sigma^2_{v^k} + \left[f\left(r\right)\right]^2\left(\sigma^2_{u^k} + \sigma^2_{\zeta^k} + 2\sigma_{\zeta^k,u^k}\right) \end{aligned} \tag{C.23}$$

Contemporaneous Parent Covariance

$$Cov\left(e^p_{f,t},n^p_{f,t}\right) = \sigma_{\bar{e}^p,\bar{n}^p}+\sigma_{\zeta^p,u^p} \tag{C.24}$$

$$Cov\left(e^p_{ft},c^p_{ft}\right) = \sigma^2_{\bar{e}^p}+\sigma_{\bar{q}^p,\bar{e}^p}+\sigma_{\bar{e}^p,\bar{n}^p}+f\left(r\right)\left(\sigma^2_{\zeta^p}+\sigma_{\zeta^p,u^p}\right) \tag{C.25}$$

$$Cov\left(n^p_{f,t},c^p_{f,t}\right) = \sigma^2_{\bar{n}^p}+\sigma_{\bar{q}^p,\bar{n}^p}+\sigma_{\bar{e}^p,\bar{n}^p}+f\left(r\right)\left(\sigma^2_{u^p}+\sigma_{\zeta^p,u^p}\right) \tag{C.26}$$

Contemporaneous Child Covariance

$$Cov\left(e^k_{f,t},n^k_{f,t}\right) = \left(\rho\gamma+\theta\lambda\right)\sigma_{\bar{e}^p,\bar{n}^p}+\gamma\lambda\sigma^2_{\bar{e}^p}+\rho\theta\sigma^2_{\bar{n}^p}+\sigma_{\delta^k,\varepsilon^k}+\sigma_{\zeta^k,u^k} \tag{C.27}$$

$$\begin{aligned} Cov\left(e^k_{f,t},c^k_{f,t}\right) &= \gamma\left(\gamma+\lambda\right)\sigma^2_{\bar{e}^p}+\theta\left(\theta+\rho\right)\sigma^2_{\bar{n}^p}+\phi\gamma\sigma_{\bar{q}^p,\bar{e}^p}+\phi\theta\sigma_{\bar{q}^p,\bar{n}^p} \\ &+ \left[\gamma\left(\rho+\theta\right)+\theta\left(\gamma+\lambda\right)\right]\sigma_{\bar{e}^p,\bar{n}^p} \\ &+ \sigma^2_{\delta^k}+\sigma_{\psi^k,\delta^k}+\sigma_{\delta^k,\varepsilon^k}+f\left(r\right)\left(\sigma^2_{\zeta^k}+\sigma_{\zeta^k,u^k}\right) \end{aligned} \tag{C.28}$$

$$\begin{aligned} Cov\left(n^k_{f,t},c^k_{f,t}\right) &= \lambda\left(\gamma+\lambda\right)\sigma^2_{\bar{e}^p}+\rho\left(\theta+\rho\right)\sigma^2_{\bar{n}^p}+\phi\lambda\sigma_{\bar{q}^p,\bar{e}^p}+\phi\rho\sigma_{\bar{q}^p,\bar{n}^p} \\ &+ \left[\lambda\left(\rho+\theta\right)+\rho\left(\gamma+\lambda\right)\right]\sigma_{\bar{e}^p,\bar{n}^p} \\ &+ \sigma^2_{\varepsilon^k}+\sigma_{\psi^k,\delta^k}+\sigma_{\psi^k,\varepsilon^k}+f\left(r\right)\left(\sigma^2_{u^k}+\sigma_{\zeta^k,u^k}\right) \end{aligned} \tag{C.29}$$

Contemporaneous Cross-Generation Covariance

$$Cov\left(e^p_{f,t},e^k_{f,t}\right) = \gamma\sigma^2_{\bar{e}^p}+\theta\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.30}$$

$$Cov\left(n^p_{f,t},n^k_{f,t}\right) = \rho\sigma^2_{\bar{n}^p}+\lambda\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.31}$$

$$\begin{aligned} Cov\left(c^p_{f,t},c^k_{f,t}\right) &= \phi\left(\sigma^2_{\bar{q}^p}+\sigma_{\bar{q}^p,\bar{e}^p}+\sigma_{\bar{q}^p,\bar{n}^p}\right)+\left(\gamma+\lambda\right)\left(\sigma^2_{\bar{e}^p}+\sigma_{\bar{q}^p,\bar{e}^p}+\sigma_{\bar{e}^p,\bar{n}^p}\right) \\ &+ \left(\rho+\theta\right)\left(\sigma^2_{\bar{n}^p}+\sigma_{\bar{q}^p,\bar{n}^p}+\sigma_{\bar{e}^p,\bar{n}^p}\right) \end{aligned} \tag{C.32}$$

$$Cov\left(e^p_{f,t},n^k_{f,t}\right) = \rho\sigma_{\bar{e}^p,\bar{n}^p}+\lambda\sigma^2_{\bar{e}^p} \tag{C.33}$$

$$Cov\left(e^p_{f,t},c^k_{f,t}\right) = \left(\gamma+\lambda\right)\sigma^2_{\bar{e}^p}+\phi\sigma_{\bar{q}^p,\bar{e}^p}+\left(\rho+\theta\right)\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.34}$$

$$Cov\left(n^p_{f,t},e^k_{f,t}\right) = \gamma\sigma_{\bar{e}^p,\bar{n}^p}+\theta\sigma^2_{\bar{n}^p} \tag{C.35}$$

$$Cov\left(n^p_{f,t},c^k_{f,t}\right) = \left(\rho+\theta\right)\sigma^2_{\bar{n}^p}+\phi\sigma_{\bar{q}^p,\bar{n}^p}+\left(\gamma+\lambda\right)\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.36}$$

$$Cov\left(c^p_{f,t},e^k_{f,t}\right) = \gamma\left(\sigma^2_{\bar{e}^p}+\sigma_{\bar{q}^p,\bar{e}^p}+\sigma_{\bar{e}^p,\bar{n}^p}\right)+\theta\left(\sigma^2_{\bar{n}^p}+\sigma_{\bar{q}^p,\bar{n}^p}+\sigma_{\bar{e}^p,\bar{n}^p}\right) \tag{C.37}$$

$$Cov\left(c^p_{f,t},n^k_{f,t}\right) = \lambda\left(\sigma^2_{\bar{e}^p}+\sigma_{\bar{q}^p,\bar{e}^p}+\sigma_{\bar{e}^p,\bar{n}^p}\right)+\rho\left(\sigma^2_{\bar{n}^p}+\sigma_{\bar{q}^p,\bar{n}^p}+\sigma_{\bar{e}^p,\bar{n}^p}\right) \tag{C.38}$$

Non-contemporaneous Covariances (lag 1) for Parent

$$Cov\left(e^p_{f,t}, e^p_{f,t+1}\right) = \sigma^2_{\bar{e}^p} \tag{C.39}$$

$$Cov\left(n^p_{f,t}, n^p_{f,t+1}\right) = \sigma^2_{\bar{n}^p} \tag{C.40}$$

$$Cov\left(c^p_{f,t}, c^p_{f,t+1}\right) = \sigma^2_{\bar{q}^p} + \sigma^2_{\bar{e}^p} + \sigma^2_{\bar{n}^p} + 2\left(\sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) \tag{C.41}$$

$$Cov\left(e^p_{f,t}, n^p_{f,t+1}\right) = \sigma_{\bar{e}^p,\bar{n}^p} \tag{C.42}$$

$$Cov\left(e^p_{f,t}, c^p_{f,t+1}\right) = \sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p} \tag{C.43}$$

$$Cov\left(n^p_{f,t}, e^p_{f,t+1}\right) = \sigma_{\bar{e}^p,\bar{n}^p} \tag{C.44}$$

$$Cov\left(n^p_{f,t}, c^p_{f,t+1}\right) = \sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p} \tag{C.45}$$

$$Cov\left(c^p_{f,t}, e^p_{f,t+1}\right) = \sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p} \tag{C.46}$$

$$Cov\left(c^p_{f,t}, n^p_{f,t+1}\right) = \sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p} \tag{C.47}$$

Non-contemporaneous Covariances (lag 1) for Child

$$Cov\left(e^k_{f,t}, e^k_{f,t+1}\right) = \gamma^2\sigma^2_{\bar{e}^p} + \theta^2\sigma^2_{\bar{n}^p} + 2\gamma\theta\sigma_{\bar{e}^p,\bar{n}^p} + \sigma^2_{\delta^k} \tag{C.48}$$

$$Cov\left(n^k_{f,t}, n^k_{f,t+1}\right) = \rho^2\sigma^2_{\bar{n}^p} + \lambda^2\sigma^2_{\bar{e}^p} + 2\rho\lambda\sigma_{\bar{e}^p,\bar{n}^p} + \sigma^2_{\varepsilon^k} \tag{C.49}$$

$$\begin{aligned} Cov\left(c^k_{f,t}, c^k_{f,t+1}\right) = {} & \phi^2\sigma^2_{\bar{q}^p} + (\gamma+\lambda)^2\sigma^2_{\bar{e}^p} + (\rho+\theta)^2\sigma^2_{\bar{n}^p} + \sigma^2_{\varepsilon^k} + \sigma^2_{\psi^k} + \sigma^2_{\delta^k} \\ & + 2\left[(\gamma+\lambda)\phi\sigma_{\bar{q}^p,\bar{e}^p} + (\rho+\theta)\phi\sigma_{\bar{q}^p,\bar{n}^p} + (\rho+\theta)(\gamma+\lambda)\sigma_{\bar{e}^p,\bar{n}^p}\right] \\ & + 2\left(\sigma_{\psi^k,\varepsilon^k} + \sigma_{\psi^k,\delta^k} + \sigma_{\delta^k,\varepsilon^k}\right) \end{aligned} \tag{C.50}$$

$$Cov\left(e^k_{f,t}, n^k_{f,t+1}\right) = (\rho\gamma+\theta\lambda)\sigma_{\bar{e}^p,\bar{n}^p} + \gamma\lambda\sigma^2_{\bar{e}^p} + \theta\rho\sigma^2_{\bar{n}^p} + \sigma_{\delta^k,\varepsilon^k} \tag{C.51}$$

$$\begin{aligned} Cov\left(e^k_{f,t}, c^k_{f,t+1}\right) = {} & \gamma\left[(\gamma+\lambda)\sigma^2_{\bar{e}^p} + (\theta+\rho)\sigma_{\bar{e}^p,\bar{n}^p} + \phi\sigma_{\bar{q}^p,\bar{e}^p}\right] + \sigma^2_{\delta^k} + \sigma_{\delta^k,\varepsilon^k} + \sigma_{\psi^k,\delta^k} \\ & + \theta\left[(\theta+\rho)\sigma^2_{\bar{n}^p} + (\gamma+\lambda)\sigma_{\bar{e}^p,\bar{n}^p} + \phi\sigma_{\bar{q}^p,\bar{n}^p}\right] \end{aligned} \tag{C.52}$$

$$Cov\left(n^k_{f,t}, e^k_{f,t+1}\right) = (\rho\gamma+\theta\lambda)\sigma_{\bar{e}^p,\bar{n}^p} + \gamma\lambda\sigma^2_{\bar{e}^p} + \theta\rho\sigma^2_{\bar{n}^p} + \sigma_{\delta^k,\varepsilon^k} \tag{C.53}$$

$$\begin{aligned} Cov\left(n^k_{f,t}, c^k_{f,t+1}\right) = {} & \rho\left[(\gamma+\lambda)\sigma^2_{\bar{n}^p} + \phi\sigma_{\bar{q}^p,\bar{n}^p} + (\gamma+\lambda)\sigma_{\bar{e}^p,\bar{n}^p}\right] + \sigma^2_{\varepsilon^k} + \sigma_{\psi^k,\varepsilon^k} + \sigma_{\delta^k,\varepsilon^k} \\ & + \lambda\left[(\gamma+\lambda)\sigma^2_{\bar{e}^p} + \phi\sigma_{\bar{q}^p,\bar{e}^p} + (\theta+\rho)\sigma_{\bar{e}^p,\bar{n}^p}\right] \end{aligned} \tag{C.54}$$

$$\begin{aligned} Cov\left(c^k_{f,t}, e^k_{f,t+1}\right) = {} & \gamma\left[(\gamma+\lambda)\sigma^2_{\bar{e}^p} + (\theta+\rho)\sigma_{\bar{e}^p,\bar{n}^p} + \phi\sigma_{\bar{q}^p,\bar{e}^p}\right] + \sigma^2_{\delta^k} + \sigma_{\delta^k,\varepsilon^k} + \sigma_{\psi^k,\delta^k} \\ & + \theta\left[(\theta+\rho)\sigma^2_{\bar{n}^p} + (\gamma+\lambda)\sigma_{\bar{e}^p,\bar{n}^p} + \phi\sigma_{\bar{q}^p,\bar{n}^p}\right] \end{aligned} \tag{C.55}$$

$$\begin{aligned} Cov\left(c^k_{f,t}, n^k_{f,t+1}\right) = {} & \rho\left[(\gamma+\lambda)\sigma^2_{\bar{n}^p} + \phi\sigma_{\bar{q}^p,\bar{n}^p} + (\gamma+\lambda)\sigma_{\bar{e}^p,\bar{n}^p}\right] + \sigma^2_{\varepsilon^k} + \sigma_{\psi^k,\varepsilon^k} + \sigma_{\delta^k,\varepsilon^k} \\ & + \lambda\left[(\gamma+\lambda)\sigma^2_{\bar{e}^p} + \phi\sigma_{\bar{q}^p,\bar{e}^p} + (\theta+\rho)\sigma_{\bar{e}^p,\bar{n}^p}\right] \end{aligned} \tag{C.56}$$

Cross-Generation Covariances: Parent at t & child at $t+1$

$$Cov\left(e^p_{f,t}, e^k_{f,t+1}\right) = \gamma\sigma^2_{\bar{e}^p} + \theta\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.57}$$

$$Cov\left(e^p_{f,t}, n^k_{f,t+1}\right) = \rho\sigma_{\bar{e}^p,\bar{n}^p} + \lambda\sigma^2_{\bar{e}^p} \tag{C.58}$$

$$Cov\left(e^p_{f,t}, c^k_{f,t+1}\right) = (\gamma+\lambda)\,\sigma^2_{\bar{e}^p} + \phi\sigma_{\bar{q}^p,\bar{e}^p} + (\rho+\theta)\,\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.59}$$

$$Cov\left(n^p_{f,t}, e^k_{f,t+1}\right) = \gamma\sigma_{\bar{e}^p,\bar{n}^p} + \theta\sigma^2_{\bar{n}^p} \tag{C.60}$$

$$Cov\left(n^p_{f,t}, n^k_{f,t+1}\right) = \rho\sigma^2_{\bar{n}^p} + \lambda\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.61}$$

$$Cov\left(n^p_{f,t}, c^k_{f,t+1}\right) = (\rho+\theta)\,\sigma^2_{\bar{n}^p} + \phi\sigma_{\bar{q}^p,\bar{n}^p} + (\gamma+\lambda)\,\sigma_{\bar{e}^p,\bar{n}^p} \tag{C.62}$$

$$Cov\left(c^p_{f,t}, e^k_{f,t+1}\right) = \gamma\left(\sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) + \theta\left(\sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) \tag{C.63}$$

$$Cov\left(c^p_{f,t}, n^k_{f,t+1}\right) = \lambda\left(\sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) + \rho\left(\sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) \tag{C.64}$$

$$\begin{aligned} Cov\left(c^p_{f,t}, c^k_{f,t+1}\right) &= \phi\left(\sigma^2_{\bar{q}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{q}^p,\bar{n}^p}\right) + (\gamma+\lambda)\left(\sigma^2_{\bar{e}^p} + \sigma_{\bar{q}^p,\bar{e}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) \\ &+ (\rho+\theta)\left(\sigma^2_{\bar{n}^p} + \sigma_{\bar{q}^p,\bar{n}^p} + \sigma_{\bar{e}^p,\bar{n}^p}\right) \end{aligned} \tag{C.65}$$

Table C.15: Transitory Shocks Estimates

	Parameters	**Estimates** (1)
Parental Transitory Shocks		.
Earnings	$\sigma^2_{\zeta^p}$	0.095 (0.006)
Other Income	$\sigma^2_{u^p}$	0.393 (0.025)
Consumption	$\sigma^2_{\upsilon^p}$	0.069 (0.004)
Earnings on Other Income	σ_{u^p,ζ^p}	-0.022 (0.004)
		.
Child Transitory Shocks		.
Earnings	$\sigma^2_{\zeta^k}$	0.097 (0.006)
Other Income	$\sigma^2_{u^k}$	0.366 (0.029)
Consumption	$\sigma^2_{\upsilon^k}$	0.086 (0.006)
Earnings on Other Income	σ_{u^k,ζ^k}	0.004 (0.007)

Note: Bootstrap standard errors (100 repetitions) in parentheses. This table uses the same sample and model specification as column (5) of Tables 4.11 and 4.12.

Identification There are 25 parameters to be identified from 48 equations. We will proceed with the identification argument in the following three steps:

(i) First, we identify 10 parameters linked to earnings, income and consumption processes for parents. Equations (C.39), (C.40), (C.18), (C.19), (C.42), (C.45), (C.46), (C.41), (C.24) and (C.20) can be considered sequentially to identify $\sigma^2_{\bar{e}^p}$, $\sigma^2_{\bar{n}^p}$, $\sigma^2_{\zeta^p}$, $\sigma^2_{u^p}$, $\sigma_{\bar{e}^p,\bar{n}^p}$, $\sigma_{\bar{q}^p,\bar{n}^p}$, $\sigma_{\bar{q}^p,\bar{e}^p}$, $\sigma^2_{\bar{q}^p}$, σ_{ζ^p,u^p} and $\sigma^2_{v^p}$ respectively.

(ii) Next, we identify 5 parameters which denote intergenerational elasticities. Equations (C.30) and (C.37) can simultaneously identify γ and θ, while ρ and λ are identified from equations (C.31) and (C.38). Finally, ϕ is identified from equation (C.32).

(iii) Lastly, the 10 parameters associated with the child's earnings, income and consumption processes are identified. Equations (C.48), (C.49), (C.21), (C.22), (C.51), (C.54), (C.55), (C.50), (C.27) and (C.23) can be considered sequentially to identify $\sigma^2_{\delta^k}$, $\sigma^2_{\varepsilon^k}$, $\sigma^2_{\zeta^k}$, $\sigma^2_{u^k}$, $\sigma_{\delta^k,\varepsilon^k}$, $\sigma_{\psi^k,\varepsilon^k}$, σ_{ψ^k,δ^k}, $\sigma^2_{\psi^k}$, σ_{ζ^k,u^k} and $\sigma^2_{v^k}$ respectively.

C.8 Random Walk Model

In this appendix, we posit an alternative to our baseline model of intergenerational persistence in individual fixed effects of income and consumption levels. We assume that the permanent component of both head earnings and other income of the family is a random walk process, and explore the extent of intergenerational persistence in the permanent innovations to these random walk components. Identification of intergenerational persistence in permanent life-cycle shocks involves calculating the growth rates of the outcome variables, which in turn implies that one can no longer identify the persistence in fixed effects, which are differenced out in growth rates.

Under this alternative view of intergenerational persistence, the model equations describing earnings and other income are:

$$e^p_{f,t} = \bar{e}^p_f + P^p_{f,t} + u^p_{f,t} \tag{C.66}$$

$$P^p_{f,t} = P^p_{f,t-1} + v^p_{f,t} \tag{C.67}$$

$$n^p_{f,t} = \bar{n}^p_f + Q^p_{f,t} + \zeta^p_{f,t} \tag{C.68}$$

$$Q^p_{f,t} = Q^p_{f,t-1} + \nu^p_{f,t} \tag{C.69}$$

A similar set of equations for earnings and other income holds true for the children. In addition, we assume that intergenerational linkages follow:

$$v^k_{f,t} = \rho v^p_{f,t} + \varepsilon^k_{f,t}$$

and

$$\nu^k_{f,t} = \lambda \nu^p_{f,t} + \theta^k_{f,t}.$$

Time differencing the income equations over successive sample years delivers the following equations:[99]

$$\begin{aligned}
\Delta_2 e^p_{f,t} &= \left(v^p_{f,t} + v^p_{f,t-1}\right) + \Delta_2 u^p_{f,t} && \text{(C.70)}\\
\Delta_2 n^p_{f,t} &= \left(\nu^p_{f,t} + \nu^p_{f,t-1}\right) + \Delta_2 \zeta^p_{f,t} && \text{(C.71)}\\
\Delta_2 e^k_{f,t} &= \rho\left(v^p_{f,t} + v^p_{f,t-1}\right) + \left(\varepsilon^k_{f,t} + \varepsilon^k_{f,t}\right) + \Delta_2 u^k_{f,t} && \text{(C.72)}\\
\Delta_2 n^k_{f,t} &= \lambda\left(\nu^p_{f,t} + \nu^p_{f,t-1}\right) + \left(\theta^k_{f,t} + \theta^k_{f,t}\right) + \Delta_2 \zeta^k_{f,t} && \text{(C.73)}
\end{aligned}$$

In this setting, the growth rate of consumption depends on the transitory and permanent innovations to earnings and other income, as well as on consumption-specific transitory heterogeneity, just as in the well-known work of Blundell, Pistaferri, and Preston (2008):

$$\Delta c^j_{f,t} = \phi_{e^j} v^j_{f,t} + \psi_{e^j} u^j_{f,t} + \psi_{n^j} \nu^j_{f,t} + \psi_{n^j} \zeta^j_{f,t} + \xi^j_{f,t} \quad where\ j = \{p,k\}.$$

The loading parameters of permanent innovations to earnings and other income in the consumption growth equation are interpreted as inverse measures of consumption insurance. For example, when ϕ_{e^j} is close to zero, permanent shocks to earnings have little or no effect on expenditure growth, which suggests the presence of effective consumption smoothing mechanisms. On the other hand, if ϕ_{e^j} is close to unity there is little insurance against innovations to permanent earnings. We also allow for the possibility of direct persistence in consumption growth so that $\xi^k_{f,t} = \gamma \xi^p_{f,t} + \chi^k_{f,t}$. This alternative model results in equations:

$$\begin{aligned}
\Delta_2 c^p_{f,t} &= \phi_{e^p}\left(v^p_{f,t} + v^p_{f,t-1}\right) + \phi_{n^p}\left(\nu^p_{f,t} + \nu^p_{f,t-1}\right)\\
&+ \psi_{e^p}\left(u^p_{f,t} + u^p_{f,t-1}\right) + \psi_{n^p}\left(\zeta^p_{f,t} + \zeta^p_{f,t-1}\right) + \left(\xi^p_{f,t} + \xi^p_{f,t-1}\right) && \text{(C.74)}
\end{aligned}$$

$$\begin{aligned}
\Delta_2 c^k_{f,t} &= \phi_{e^k}\left[\rho\left(v^p_{f,t} + v^p_{f,t-1}\right) + \left(\varepsilon^k_{f,t} + \varepsilon^k_{f,t-1}\right)\right] + \psi_{e^k}\left(u^k_{f,t} + u^k_{f,t-1}\right)\\
&+ \phi_{n^k}\left[\lambda\left(\nu^p_{f,t} + \nu^p_{f,t-1}\right) + \left(\theta^k_{f,t} + \theta^k_{f,t-1}\right)\right] + \psi_{n^k}\left(\zeta^k_{f,t} + \zeta^k_{f,t-1}\right)\\
&+ \gamma\left(\xi^p_{f,t} + \xi^p_{f,t-1}\right) + \left(\chi^k_{f,t} + \chi^k_{f,t-1}\right) && \text{(C.75)}
\end{aligned}$$

[99]In equations (C.70) through (C.75), we use the notation $\Delta_2 x_t \equiv x_t - x_{t-2}$ to denote the two-year time difference for any variable x_t. Since PSID data are only available every two years after 1998, we consider two-year time differences throughout so as to use data from both pre and post 1998 interview rounds.

C.8.1 Moment Conditions

Parent Variance

$$Var\left(\Delta_2 e^p_{f,t}\right) = 2\left(\sigma^2_{\upsilon^p} + \sigma^2_{u^p}\right) \tag{C.76}$$

$$Var\left(\Delta_2 n^p_{f,t}\right) = 2\left(\sigma^2_{\nu^p} + \sigma^2_{\zeta^p}\right) \tag{C.77}$$

$$Var\left(\Delta_2 c^p_{f,t}\right) = 2\left(\phi^2_{e^p}\sigma^2_{\upsilon^p} + \phi^2_{n^p}\sigma^2_{\nu^p} + \psi^2_{e^p}\sigma^2_{u^p} + \psi^2_{n^p}\sigma^2_{\zeta^p} + \sigma^2_{\xi^p}\right) \tag{C.78}$$

Child Variance

$$Var\left(\Delta_2 e^k_{f,t}\right) = 2\left(\rho^2\sigma^2_{\upsilon^p} + \sigma^2_{u^k} + \sigma^2_{\varepsilon^k}\right) \tag{C.79}$$

$$Var\left(\Delta_2 n^k_{f,t}\right) = 2\left(\lambda^2\sigma^2_{\nu^p} + \sigma^2_{\zeta^k} + \sigma^2_{\theta^k}\right) \tag{C.80}$$

$$\begin{aligned} Var\left(\Delta_2 c^k_{f,t}\right) &= 2\left(\rho^2\phi^2_{e^k}\sigma^2_{\upsilon^p} + \phi^2_{e^k}\sigma^2_{\varepsilon^k} + \psi^2_{e^k}\sigma^2_{u^k}\right) \\ &+ 2\left(\lambda^2\phi^2_{n^k}\sigma^2_{\nu^p} + \phi^2_{n^k}\sigma^2_{\theta^k} + \psi^2_{n^k}\sigma^2_{\zeta^k} + \gamma^2\sigma^2_{\xi^p} + \sigma^2_{\chi^k}\right) \end{aligned} \tag{C.81}$$

Contemporaneous Parent Covariance

$$Cov\left(\Delta_2 e^p_{f,t}, \Delta_2 c^p_{f,t}\right) = 2\phi_{e^p}\sigma^2_{\upsilon^p} + \psi_{e^p}\sigma^2_{u^p} \tag{C.82}$$

$$Cov\left(\Delta_2 n^p_{f,t}, \Delta_2 c^p_{f,t}\right) = 2\phi_{n^p}\sigma^2_{\nu^p} + \psi_{n^p}\sigma^2_{\zeta^p} \tag{C.83}$$

Contemporaneous Child Covariance

$$Cov\left(\Delta_2 e^k_{f,t}, \Delta_2 c^k_{f,t}\right) = 2\rho^2\phi_{e^k}\sigma^2_{\upsilon^p} + 2\phi_{e^k}\sigma^2_{\varepsilon^k} + \psi_{e^k}\sigma^2_{u^k} \tag{C.84}$$

$$Cov\left(\Delta_2 n^k_{f,t}, \Delta_2 c^k_{f,t}\right) = 2\lambda^2\phi_{n^k}\sigma^2_{\nu^p} + 2\phi_{n^k}\sigma^2_{\theta^k} + \psi_{n^k}\sigma^2_{\zeta^k} \tag{C.85}$$

Contemporaneous Cross-Generation Covariance

$$Cov\left(\Delta_2 e^p_{f,t}, \Delta_2 e^k_{f,t}\right) = 2\rho\sigma^2_{\upsilon^p} \tag{C.86}$$

$$Cov\left(\Delta_2 n^p_{f,t}, \Delta_2 n^k_{f,t}\right) = 2\lambda\sigma^2_{\nu^p} \tag{C.87}$$

$$Cov\left(\Delta_2 c^p_{f,t}, \Delta_2 c^k_{f,t}\right) = 2\left(\rho\phi_{e^p}\phi_{e^k}\sigma^2_{\upsilon^p} + \lambda\phi_{n^p}\phi_{n^k}\sigma^2_{\nu^p} + \gamma\sigma^2_{\xi^p}\right) \tag{C.88}$$

$$Cov\left(\Delta_2 e^p_{f,t}, \Delta_2 c^k_{f,t}\right) = 2\rho\phi_{e^k}\sigma^2_{\upsilon^p} \tag{C.89}$$

$$Cov\left(\Delta_2 n^p_{f,t}, \Delta_2 c^k_{f,t}\right) = 2\lambda\phi_{n^k}\sigma^2_{\nu^p} \tag{C.90}$$

$$Cov\left(\Delta_2 c^p_{f,t}, \Delta_2 e^k_{f,t}\right) = 2\rho\phi_{e^p}\sigma^2_{\upsilon^p} \tag{C.91}$$

$$Cov\left(\Delta_2 c^p_{f,t}, \Delta_2 n^k_{f,t}\right) = 2\lambda\phi_{n^p}\sigma^2_{\nu^p} \tag{C.92}$$

Non-contemporaneous Covariances (lag 2) for Parent

$$Cov\left(\Delta_2 e^p_{f,t}, \Delta_2 e^p_{f,t+2}\right) = -\sigma^2_{u^p} \tag{C.93}$$

$$Cov\left(\Delta_2 n^p_{f,t}, \Delta_2 n^p_{f,t+2}\right) = -\sigma^2_{\zeta^p} \tag{C.94}$$

$$Cov\left(\Delta_2 c^p_{f,t}, \Delta_2 e^p_{f,t+2}\right) = -\psi_{e^p}\sigma^2_{u^p} \tag{C.95}$$

$$Cov\left(\Delta_2 c^p_{f,t}, \Delta_2 n^p_{f,t+2}\right) = -\psi_{n^p}\sigma^2_{\zeta^p} \tag{C.96}$$

Non-contemporaneous Covariances (lag 2) for Child

$$Cov\left(\Delta_2 e^k_{f,t}, \Delta_2 e^k_{f,t+2}\right) = -\sigma^2_{u^k} \tag{C.97}$$

$$Cov\left(\Delta_2 n^k_{f,t}, \Delta_2 n^k_{f,t+2}\right) = -\sigma^2_{\zeta^k} \tag{C.98}$$

$$Cov\left(\Delta_2 c^k_{f,t}, \Delta_2 e^k_{f,t+2}\right) = -\psi_{e^k}\sigma^2_{u^k} \tag{C.99}$$

$$Cov\left(\Delta_2 c^k_{f,t}, \Delta_2 n^k_{f,t+2}\right) = -\psi_{n^k}\sigma^2_{\zeta^k} \tag{C.100}$$

C.8.2 Identification

There are 21 parameters to be identified from 25 moment conditions. It is straightforward to see the identification of $\sigma^2_{u^p}$, $\sigma^2_{\zeta^p}$, ψ_{e^p}, ψ_{n^p}, $\sigma^2_{u^k}$, $\sigma^2_{\zeta^k}$, ψ_{e^k} and ψ_{n^k} from equations (C.93) through (C.100). Subsequently, $\sigma^2_{\upsilon^p}$ and $\sigma^2_{\nu^p}$ can be identified from equations (C.76) and (C.77). This allows identification of ρ and λ from equations (C.86) and (C.87); and consequently ϕ_{e^k}, ϕ_{n^k}, ϕ_{e^p} and ϕ_{n^p} from equations (C.89) through (C.92) respectively. Now, equations (C.78), (C.79) and (C.80) can identify $\sigma^2_{\xi^p}$, $\sigma^2_{\varepsilon^k}$ and $\sigma^2_{\theta^k}$ respectively. Finally, γ is identified from equation (C.88), which leaves $\sigma^2_{\chi^k}$ to be identified from (C.81).

C.8.3 Estimates

In Table C.16, we present two sets of estimates for this random walk model. The first set is based on imputed expenditure data; the second set is obtained using only directly observed food expenditures as a measure of consumption. Blundell, Pistaferri, and Preston (2008) point out that *"...using food would provide an estimate of insurance that is ...higher than with imputed consumption data"* and *"...may give misleading evidence on the size and the stability of the insurance parameters."* Not surprisingly, therefore, Table C.17 shows that we estimate higher value of consumption insurance when using food expenditures rather than imputed consumption data. Table C.16 shows that innovations to earnings, other income and consumption display no statistically significant persistence across generations. With the caveat that first-differenced data can exacerbate measurement error and reduce significance, we find no evidence of intergenerational linkages in the accrual rate of permanent innovations over the life-cycle.

Table C.16: Intergenerational Growth Elasticities

	Parameters	**Imputed** (1)	**Food** (2)
Earnings Growth	ρ	0.241	0.256
		(0.161)	(0.193)
Other Income Growth	λ	0.094	0.095
		(0.071)	(0.059)
Consumption Growth Shifter	γ	0.009	0.047
		(0.048)	(0.056)
No. of Parent-Child Pairs	N	760	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. Year and cohort effects have been removed.

Table C.17: Partial Insurance Parameters

	Parameters	**Imputed** (1)	**Food** (2)
Parents			
Permanent Earnings	ϕ_e^p	0.230	0.104
		(0.037)	(0.085)
Permanent Other Income	ϕ_n^p	0.069	0.033
		(0.017)	(0.025)
Transitory Earnings	ψ_e^p	0.147	0.057
		(0.034)	(0.094)
Transitory Other Income	ψ_n^p	0.033	-0.047
		(0.042)	(0.066)
Children			
Permanent Earnings	ϕ_e^k	0.237	0.034
		(0.053)	(0.102)
Permanent Other Income	ϕ_n^k	0.127	0.076
		(0.021)	(0.022)
Transitory Earnings	ψ_e^k	0.201	0.023
		(0.036)	(0.067)
Transitory Other Income	ψ_n^k	0.046	-0.042
		(0.025)	(0.065)
No. of Parent-Child Pairs	N	760	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. Data are purged of year and cohort effects.

Table C.18: Variances of Shocks

	Parameters	Imputed (1)	Food (2)
Parental Shocks			
Transitory Earnings	$\sigma^2_{u^p}$	0.048	0.048
		(0.005)	(0.004)
Transitory Other Income	$\sigma^2_{\zeta^p}$	0.068	0.068
		(0.015)	(0.016)
Permanent Earnings	$\sigma^2_{\upsilon^p}$	0.033	0.033
		(0.004)	(0.004)
Permanent Other Income	$\sigma^2_{\nu^p}$	0.108	0.107
		(0.012)	(0.013)
Consumption Growth	$\sigma^2_{\xi^p}$	0.017	0.070
		(0.001)	(0.004)
Child Shocks			
Transitory Earnings	$\sigma^2_{u^k}$	0.048	0.049
		(0.005)	(0.006)
Transitory Other Income	$\sigma^2_{\zeta^k}$	0.087	0.087
		(0.013)	(0.013)
Permanent Earnings	$\sigma^2_{\varepsilon^k}$	0.024	0.023
		(0.004)	(0.005)
Permanent Other Income	$\sigma^2_{\theta^k}$	0.095	0.095
		(0.014)	(0.015)
Consumption Growth	$\sigma^2_{\chi^k}$	0.016	0.088
		(0.001)	(0.006)
No. of Parent-Child Pairs	N	760	760

Note: Bootstrap standard errors (100 repetitions) in parentheses. Data are purged of year and cohort effects.

Table C.19: Growth Model Moments

Moments	**Imputed** (1)	**Food** (2)
$Var\left(\Delta_2 e^p_{f,t}\right)$	0.161	0.161
	(0.009)	(0.007)
$Var\left(\Delta_2 n^p_{f,t}\right)$	0.351	0.351
	(0.036)	(0.036)
$Var\left(\Delta_2 c^p_{f,t}\right)$	0.041	0.142
	(0.002)	(0.007)
$Var\left(\Delta_2 e^k_{f,t}\right)$	0.148	0.148
	(0.01)	(0.009)
$Var\left(\Delta_2 n^k_{f,t}\right)$	0.366	0.366
	(0.033)	(0.034)
$Var\left(\Delta_2 c^k_{f,t}\right)$	0.042	0.177
	(0.001)	(0.011)
$\text{Cov}\left(\Delta_2 e^p_{f,t} \Delta_2 e^k_{f,t}\right)$	0.017	0.017
	(0.011)	(0.012)
$\text{Cov}\left(\Delta_2 n^p_{f,t} \Delta_2 n^k_{f,t}\right)$	0.020	0.020
	(0.014)	(0.013)
$\text{Cov}\left(\Delta_2 c^p_{f,t} \Delta_2 c^k_{f,t}\right)$	0.001	0.007
	(0.002)	(0.008)
$\text{Cov}\left(\Delta_2 e^p_{f,t} \Delta_2 e^p_{f,t+2}\right)$	-0.048	-0.048
	(0.005)	(0.004)
$\text{Cov}\left(\Delta_2 n^p_{f,t} \Delta_2 n^p_{f,t+2}\right)$	-0.068	-0.068
	(0.015)	(0.016)
$\text{Cov}\left(\Delta_2 e^k_{f,t} \Delta_2 e^k_{f,t+2}\right)$	-0.049	-0.049
	(0.005)	(0.006)
$\text{Cov}\left(\Delta_2 n^k_{f,t} \Delta_2 n^k_{f,t+2}\right)$	-0.087	-0.087
	(0.013)	(0.013)
$\text{Cov}\left(\Delta_2 e^p_{f,t} \Delta_2 c^p_{f,t}\right)$	0.023	0.011
	(0.002)	(0.003)
$\text{Cov}\left(\Delta_2 e^p_{f,t+2} \Delta_2 c^p_{f,t}\right)$	-0.006	-0.002
	(0.002)	(0.004)
$\text{Cov}\left(\Delta_2 n^p_{f,t} \Delta_2 c^p_{f,t}\right)$	0.017	0.004
	(0.003)	(0.003)
$\text{Cov}\left(\Delta_2 n^p_{f,t+2} \Delta_2 c^p_{f,t}\right)$	-0.002	0.003
	(0.002)	(0.005)
$\text{Cov}\left(\Delta_2 e^k_{f,t} \Delta_2 c^k_{f,t}\right)$	0.023	0.004
	(0.002)	(0.003)
$\text{Cov}\left(\Delta_2 e^k_{f,t+2} \Delta_2 c^k_{f,t}\right)$	-0.008	0.000
	(0.002)	(0.003)
$\text{Cov}\left(\Delta_2 n^k_{f,t} \Delta_2 c^k_{f,t}\right)$	0.028	0.010
	(0.003)	(0.004)
$\text{Cov}\left(\Delta_2 n^k_{f,t+2} \Delta_2 c^k_{f,t}\right)$	-0.004	0.003
	(0.002)	(0.005)
$\text{Cov}\left(\Delta_2 e^p_{f,t} \Delta_2 c^k_{f,t}\right)$	-0.001	-0.003
	(0.004)	(0.009)
$\text{Cov}\left(\Delta_2 n^p_{f,t} \Delta_2 c^k_{f,t}\right)$	0.005	0.006
	(0.003)	(0.006)
$\text{Cov}\left(\Delta_2 c^p_{f,t} \Delta_2 e^k_{f,t}\right)$	0.001	-0.003
	(0.003)	(0.006)
$\text{Cov}\left(\Delta_2 c^p_{f,t} \Delta_2 n^k_{f,t}\right)$	-0.003	-0.002
	(0.008)	(0.011)

Note: These empirical moments are used to generate the parameter estimates in Tables C.16, C.17 and C.18 through GMM. Bootstrap standard errors are reported in parentheses.