From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Economics

2014

Essays on Changing Labor Markets and the Macroeconomy

vom Lehn, John Christian

Princeton, NJ : Princeton University

http://arks.princeton.edu/ark:/88435/dsp014q77fr49v

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

ESSAYS ON CHANGING LABOR MARKETS AND THE MACROECONOMY

JOHN CHRISTIAN VOM LEHN JR.

A DISSERTATION

PRESENTED TO THE FACULTY

OF PRINCETON UNIVERSITY

IN CANDIDACY FOR THE DEGREE

OF DOCTOR OF PHILOSOPHY

RECOMMENDED FOR ACCEPTANCE

BY THE DEPARTMENT OF

ECONOMICS

ADVISER: RICHARD ROGERSON

JUNE 2014

© Copyright by John Christian vom Lehn Jr., 2014.

All rights reserved.

Abstract

This thesis furthers our understanding of several recent changes affecting labor markets and the macroeconomy. Chapter 1 addresses the role of technological change in the recent polarization of the labor market through a quantitative assessment of the relative price of investment channel, using an extension of the standard growth model, allowing for an endogenously determined supply and demand for occupational labor. The model is able to account for a substantial fraction of changes in relative employment and wages across occupations, but only through the year 2000, suggesting a reduced role for technological change in the last decade. The model is also able to capture 25% of the observed decline in the labor share and 35% of the decline in the employment to population ratio after 2000, suggesting a substantial role for secular factors in recently sluggish labor markets.

Chapter 2 addresses the recent decline of the U.S. labor share and its relation to changes across occupations specializing in different tasks, reaching two main conclusions. First, the labor share decline is linked to the decline of compensation of workers employed in routine occupations. Second, until the year 2000, most of this decline in routine compensation was offset by growth in income paid to high-skill occupations, but this growth ceased in the year 2000, generating the accelerated decline in the labor share thereafter. I show that this recent slowdown of high skill compensation is specific to certain detailed industries and occupations and suggest implications of this evidence for theories of the slowdown.

Chapter 3 makes several points regarding changes in the cyclicality of labor productivity. I show the that change in cyclical productivity is not statistically synonymous with recent jobless recoveries, but is synonymous with the change in the relative cyclical volatility of hours and output. I show there is substantial heterogeneity in this relative volatility within industries and across time, and that aggregate changes in this relative volatility are also partially explained by changes in between industry co-movements. Existing theories focusing on changes in labor frictions appear unable to fully reconcile these facts.

Acknowledgments

I am deeply indebted to my adviser, Richard Rogerson, for frequent feedback and advice throughout the process of writing this dissertation. I am also very grateful to my committee, Esteban Rossi-Hansberg and Mark Watson, for their comments, suggestions and advice on my work. I have also benefited from comments from and conversations with many faculty members during my time at Princeton, particularly Mark Aguiar, Jesus Fernandez-Villaverde, Mike Golosov, Oleg Itskhoki, Greg Kaplan, Nobu Kiyotaki, Ben Moll, Ezra Oberfield, Chris Sims, and Jon Vogel. I am also grateful to Riccardo DiCecio for making available his data on the quality-adjusted price of equipment investment.

I have greatly enjoyed the fellowship of many fellow graduate students during the course of writing this dissertation. I am particularly grateful to my good friend, Tom Winberry, for our frequent conversations and discussion. I am also very grateful for the association with and feedback I have received from Qingqing Cao, Tom Carter, Wei Cui, Diogo Guillen, Pedro Olea, Ricardo Reyes-Heroles, Sharon Traiberman, and Gabriel Tenorio. I am also appreciative of all the attendants and participants at the Princeton Macro Student Workshop.

And of course, I am forever grateful to the support and kindness of my family, particularly my wife, Brittany, who has provided a wonderful listening ear and been a constant support throughout this process.

Contents

Abstract iii

Acknowledgements iv

1 Labor Market Polarization and the Decline of Routine Work: A Quantitative Analysis 1

1.1 Introduction 1

1.2 Polarization, the Decline of Routine Work and Aggregate Labor Market Changes . 5

1.2.1 Polarization, Routine Work, and Employment Rates: Background 5

1.2.1.1 Polarization and Routine Work 5

1.2.1.2 Employment Participation and Skill Levels 9

1.2.2 Labor Market Trend Dynamics 10

1.2.2.1 Occupational Employment Shares 11

1.2.2.2 Relative Hourly Earnings Across Occupations 11

1.2.2.3 Employment to Population Ratio 12

1.2.2.4 Labor Share of Income 14

1.3 Model 18

1.3.1 Environment 18

1.3.1.1 Technology 18

1.3.1.2 Household 20

1.3.1.3 Firm 23

1.3.2 Equilibrium 23

1.4 Driving Forces and Solution Method 24

1.4.1 Measuring Driving Forces . . . 24
1.4.2 Solution Method . . . 26
1.5 Calibration . . . 27
1.5.1 Calibration of Other Parameters . . . 28
1.5.2 Calibrating Elasticity Parameters . . . 29
1.5.3 Parameter Values . . . 31
1.6 Results . . . 34
1.6.1 Initial Intuition . . . 34
1.6.2 Model Predicted Outcomes . . . 35
1.6.2.1 Employment Shares . . . 35
1.6.2.2 Relative Hourly Earnings . . . 37
1.6.2.3 Employment to Population Ratio . . . 39
1.6.2.4 Labor's Share of Income . . . 41
1.6.3 Underlying Mechanisms . . . 43
1.6.3.1 Employment . . . 43
1.6.3.2 Hourly Earnings . . . 45
1.6.4 Summary . . . 47
1.7 Conclusion . . . 49
1.8 Appendix . . . 51
1.8.1 Data Sources and Construction . . . 51
1.8.1.1 Occupational and Industry Classifications . . . 51
1.8.1.2 CPS Basic Monthly Files . . . 52
1.8.1.3 CPS March Supplement Files . . . 53
1.8.1.4 Relative Price Data . . . 55
1.8.2 Industry Robustness Checks for the Measured Decline of Routine Work . . . 55
1.8.3 Solution and Calibration Methods . . . 59
1.8.3.1 Solution Method . . . 59

1.8.3.2 Calibration Approach 60

2 Occupational Tasks and the Decline of the U.S. Labor Share 65

2.1 Introduction 65

2.2 The Decline of the Labor Share 68

2.3 Linking Labor Share Changes to Occupational Tasks 71

2.3.1 Occupational Labor Shares 71

2.3.2 Controlling for Industry Composition 74

2.3.3 Occupational Labor Share Changes Within Industries 77

2.3.4 Within Industry Relationships Between Tasks and Labor Shares 81

2.4 Understanding the Slowdown of Abstract Income 87

2.4.1 Occupational Differences Between Sectors 87

2.4.2 Hours vs. Wages in Abstract Occupational Labor Shares 92

2.4.3 Possible Explanations 93

2.5 Conclusion 98

2.6 Appendix 99

2.6.1 Decomposing Total Compensation across Occupations 99

2.6.2 Consistent Industry Codes and Detailed Abstract Occupation Codes 101

3 Acyclical Labor Productivity: It's Not as Simple as You Think 104

3.1 Introduction 104

3.2 The Emergence of Acyclical Labor Productivity 107

3.3 Acyclical Productivity and Aggregate Changes 110

3.3.1 POINT #1: Acyclical productivity is not statistically synonymous with jobless recoveries 110

3.3.2 POINT #2: Acyclical productivity is statistically synonymous with changes in relative volatilities occurring in the Great Moderation 113

3.4 Industry Level Changes 119

3.4.1 POINT #3: There is a great deal of heterogeneity in the relative volatility of hours and output across industries and across time. 121
3.4.2 POINT #4: Aggregate changes in the relative volatility of hours and output are also partially explained by changes in between industry co-movements. 125
3.5 Evaluating Existing Explanations 129
3.6 Conclusions 131
3.7 Appendix 132
3.7.1 Proof of Lemma Relating Relative Volatility and Output-Productivity Correlation 132
3.7.2 Complete Two-Sector Decomposition of Aggregate Relative Volatility . . . 134

Bibliography **139**

Chapter 1

Labor Market Polarization and the Decline of Routine Work: A Quantitative Analysis

1.1 Introduction

Much attention has been given to changes in wage inequality occurring in recent decades. A substantial literature has focused on understanding the growth in both the supply and wages of skilled workers relative to unskilled workers, and identified skill-biased technical change as a leading explanation for these changes.[1] However, as shown in Autor et al. (2006, 2008), this framework for understanding wage inequality has been unable to capture the recently observed polarization of wage and employment growth across skill levels. Since the 1980s, wage and employment growth has been biased against middle-skilled occupations, with the fastest growth occurring in high and low skill occupations. As a result, a growing literature on this polarization of the labor market has not only thought about worker skills, but the tasks they perform in different occupations.[2] A common conclusion of this literature is that middle skilled jobs tend to specialize in performing

[1] Skilled workers are generally defined as those possessing at least a college degree. See Katz and Autor (1999) for a survey of this early literature.

[2] See Autor et al. (2003) and Goos and Manning (2007) for early work relating job tasks and labor market outcomes, and see Acemoglu and Autor (2011) for a survey of the literature that has risen in its wake.

routine tasks - tasks that are repetitive, follow strict rules or procedures, and require accuracy and precision.

One prominent hypothesis for labor market polarization is the presence of routine-biased technical change, a refinement of previous theories of skill-biased technical change. This idea, originally set forth in Autor et al. (2003), conjectures that these routine tasks can also be performed by machines, particularly information and communications technologies (ICT), and that the recently progress (and declining price) of ICT technology has motivated a labor demand shift away from routine labor inputs towards other more complementary, non-routine labor inputs at both ends of the skill distribution.[3] Persuasive suggestive evidence in favor of this hypothesis has been presented recently in Goos et al. (2010), Autor and Dorn (2013) and Michaels et al. (2014).

However, this notion of routine-biased technical change has yet to be formally evaluated quantitatively. In this paper, I take a first step towards a formal quantitative evaluation of this hypothesis by developing a dynamic general equilibrium model of occupational labor supply and demand and comparing the model's predicted outcomes to the data on changes in employment and hourly earnings across occupational groups. My model augments the standard growth model by adding three key features: (1) the occupational tasks framework of Autor and Dorn (2013), (2) a Roy model over occupational choice, patterned after Jung and Mercenier (2010) and Cortes (2012), and (3) equipment investment specific technical change, as in Greenwood et al. (1997). The model allows for endogenous determination of the supply and demand for labor in occupational tasks (abstract, routine and manual) in response to exogenous technological change, occurring through the equipment investment channel, creating a rich framework to quantitatively explore the recently observed changes in occupational employment and earnings.

Further, in contrast to much of the related literature, I give particular attention to the labor supply decision of employment vs. non-employment, and my model explicitly incorporates this margin. Two observations motivate the inclusion of this margin. First, employment to population ratios across the skill distribution do not display polarization, but have declined far more for

[3]It is also possible that recent technological change has also facilitated greater offshoring of routine labor inputs. See Grossman and Rossi-Hansberg (2008) for a discussion.

lower skilled workers than middle and higher skilled workers.[4] Understanding this disconnect between employment shares by occupational skill and employment to population ratios by worker skill is important for understanding recent changes in inequality. Second, the aggregate employment to population ratio has decreased significantly since 2000 and there is an ongoing debate about whether this decline has been due to cyclical or structural factors. Inclusion of this employment decision margin allows me to consider the extent to which the structural forces leading to polarization might be affecting the aggregate employment rate.

With this framework in place, I feed in a measured series for technological change and evaluate the extent to which the model can account for: (1) changes in employment shares across occupations differentiated by their task specialties, (2) changes in relative hourly earnings among these same tasks, and (3) changes in the employment to population ratio in the aggregate and across skill levels. In performing this quantitative exercise, I give particular attention to the time paths of these changes, and not just the long run differences between two periods.

To implement this quantitative evaluation, I measure technological change using the relative price of equipment investment and consumption goods, as in Cummins and Violante (2002). Given this series, an important question for the quantitative implementation is how to obtain values for parameters governing the elasticities of substitution between different capital and labor task inputs. With no existing estimates of these exact elasticities as they appear in my model, as a starting point, I calibrate these parameters to match changes in employment and equipment quantities in the 1980s and early 1990s. By restricting the calibration time window to this period, I avoid potential conflation with other plausible driving forces of polarization that have not emerged until more recently, such as recent increases international trade, particularly with China.

I find that technological change plays a substantial role in explaining polarization and other recent labor market outcomes, but that this role may have somewhat diminished in the recent decade. In particular, my model accounts for all observed changes in occupational employment shares in the 1990s, but substantially overstates trend changes in abstract and manual occupations

[4]This fact has been shown for longer periods of time in Autor (2010), and for more recent declines in the 2000s by Moffitt (2012).

since the year 2000. The model is able to capture well the qualitative changes in relative hourly earnings, and accounts for 35-55% of the decline in routine hourly earnings relative to hourly earnings in abstract occupations. Further, the model is consistent with qualitative observations regarding changes in employment rates across skill levels and explains 35% of the decline in the employment to population ratio since 2001, suggesting a substantial role for secular factors in explaining ongoing sluggish growth of employment.

Another salient feature of recent secular changes in the US labor market is the substantial drop in labor's share of income. As routine-biased technological change has potentially strong implications for this aggregate statistic, I also consider the model's predictions for the labor share. I show that the model can capture 25% of the observed decline in the labor share, but is unable to account for the more recent sharp decline in this series since the year 2000.

This paper relates to the literature on technological change and occupational tasks as cited above, as well as recent efforts to disentangle the dual roles of trade and technology in determining labor market outcomes, as in Autor, Dorn and Hanson (2013b). Further, the emphasis on quantitatively evaluating the role of technological change in determining labor market outcomes is in the same spirit as the work done by Krusell et al. (2000), who performed a similar evaluation of the role of equipment investment specific technological change in explaining the skill premium. In addition to focusing on three types of tasks instead of just two types of skills, in contrast to Krusell et al., I do not restrict labor supply to be exogenous, instead allowing for endogenous employment and occupational choice decisions, giving even a richer and broader framework for understanding and evaluating recent changes occurring in the labor market.

The remainder of the paper is structured as follows. Section 2 gives background on labor market polarization and its connection to worker tasks, and then documents the four facts considered in the quantitative evaluation: occupational employment and hourly earnings trends in abstract, routine and manual occupations, as well as changes in employment to population ratios and the labor share. Section 3 lays out the model framework, in particular the extension of the standard growth model to include occupational tasks, a Roy model, and equipment investment specific technological change.

Section 4 outlines procedures for how I measure the model's driving forces and how I solve the model, and Section 5 presents the calibration strategy. Section 6 presents the results from feeding in the data series on the relative price of equipment investment, gives some analysis of these results and briefly discusses possible extensions. Section 7 concludes.

1.2 Polarization, the Decline of Routine Work and Aggregate Labor Market Changes

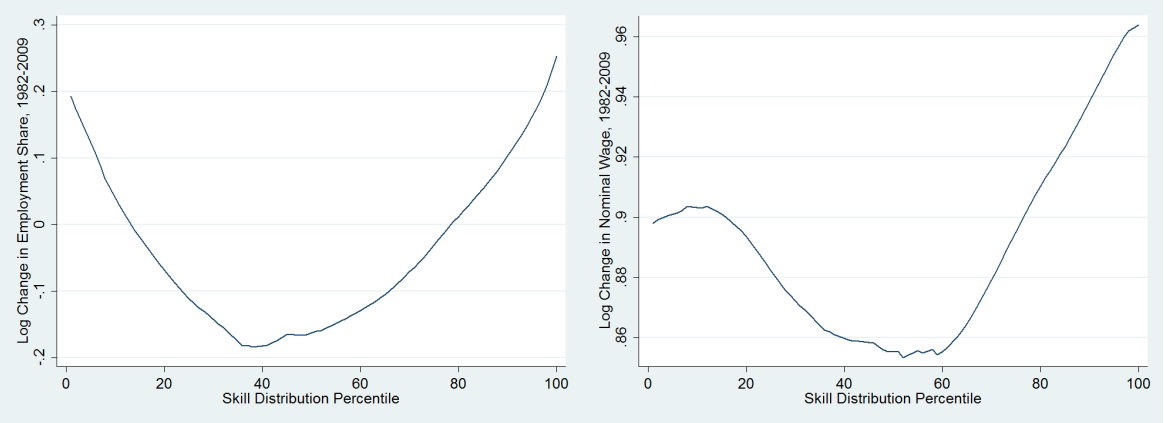
In this section, I start by giving background on the recent polarization of the labor market and its connection to occupations specializing in performing routine tasks. I also highlight how polarization has not been observed in changes in employment rates across different skill levels. Following this background, I present the four sets of facts I will focus on in my quantitative analysis: (1) occupational employment shares (2) relative hourly earnings across occupations (3) employment to population ratios, both aggregate and across skill populations (4) the labor share of income. I use data from the monthly Current Population Survey files to construct evidence on occupational employment shares and I use income data from the CPS March Supplement to construct measures of occupational hourly earnings from 1976-2011. A more detailed description of data construction used for these results and the results presented in subsequent sections is presented in the Appendix. Occupational employment and hourly earnings trends are reported for the non-farm business sector, but they are robust to using total economy data.

1.2.1 Polarization, Routine Work, and Employment Rates: Background

1.2.1.1 Polarization and Routine Work

Figure 1.1 plots the growth in wages and the growth in employment shares across the occupational skill distribution between 1982 and 2009, where the mean of occupational log wages in 1982 is

Figure 1.1: Employment and Wage Polarization Across the Skill Distribution, 1982-2009



Data comes from the March Annual Supplement to the CPS from 1982-2009 for full-time, full-year workers above the age of 16. Occupations are ranked by their log mean wage between the years 1982 and 1984, where labor supply weights are used in taking the average (the product of weeks and hours worked). With this ranking scheme, occupations are divided across 100 percentiles of the wage distribution. The left panel documents employment growth between 1982 and 2009 - the log change in the share of employment for occupations in a given skill percentile. Wage growth between 1982 and 2009 is documented in the right panel, where growth is given by log changes in the average nominal wage for a the composition of occupations in a given percentile. The smoothed lines shown above come from a locally weighted regression, with bandwidth of 0.8, of the growth rates of different percentiles, where the data has been winsorized at a 5% rate. The results, however, are robust to this procedure in dealing with outliers. The construction of these plots is very similar to those presented in Autor and Dorn (2013).

used as a proxy for skill[5]. Employment and wages have grown most rapidly for occupations in the lowest and highest skill percentiles, and have grown much slower for occupations in the roughly the middle 60% of the skill distribution. This is what has been termed the polarization of the labor market.

To better understand this phenomenon, economists have focused attention on the different types of tasks performed in occupations across the skill distribution - see Acemoglu and Autor (2011) for a recent survey. A common framework that has emerged from this literature features two broad groups of tasks performed by workers - routine tasks and non-routine tasks. A routine task is highly repetitive, follows a strict set of rules or procedures, and requires precision and accuracy.

[5]This figure follows the same general construction methods as those in the literature, particularly Acemoglu and Autor (2011) and Autor and Dorn (2013). Results are robust to excluding the recent recession and considering the growth through the year 2007. Additional details regarding data sources and construction can be found in the Appendix.

On the other hand, non-routine tasks require either creativity, abstract reasoning, persuasion or some degree of physical exertion and dexterity or personal interaction. While different studies have further decomposed these two broad groups of tasks in different ways, I follow Autor and Dorn (2013) and consider three types of tasks - abstract, routine and manual tasks.[6] Abstract tasks are those non-routine tasks which require intuition, creativity, persuasion and/or abstract reasoning and can be thought of high-skilled non-routine tasks. Manual tasks are those non-routine tasks which require physical dexterity, strength and/or interpersonal interaction and are thought of as low-skilled non-routine tasks.

Although almost all jobs perform each of these three types of tasks to some extent, occupations can be sorted by the tasks which are most prominent in the work they do. Autor and Dorn (2013) using data from the Dictionary of Occupational Titles to identify different task intensities across occupations and classify occupations according to these three task groups; I use these same classifications. Examples of occupations that specialize in each of these three types of tasks are:

Abstract: teachers, doctors, lawyers, engineers, economists

Routine: secretaries, bookkeepers, retail salespersons, machine operators, assembly line workers

Manual: hairdressers, auto mechanics, bus drivers, housekeepers, construction laborers

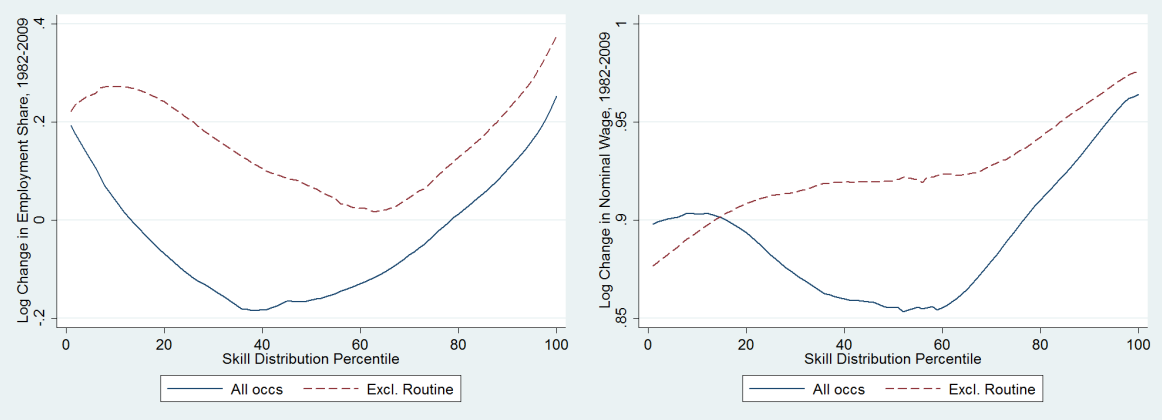
Autor and Dorn (2013) have shown that these three occupations map fairly well into the wage and educational distribution. Abstract occupations tend to be high-skilled, high-paying jobs, routine occupations more middle-educated and middle-paying, and manual occupations are generally lower-skilled, lower-paying jobs.

As routine jobs tend to be middle-skilled jobs, they are closely related to labor market polarization. To highlight this, I consider growth in employment and wages across the skill distribution when routine occupations are omitted.[7] Figure 1.2 shows these results. With routine occupations

[6]The seminal work on tasks and labor market outcomes by Autor et al. (2003) originally introduced five categories of occupations; more recent work by the authors has focused on three or four categories, as in Autor and Dorn (2013) and Acemoglu and Autor (2011).

[7]In removing routine occupations, they are only removed from calculating the growth of wages and employment in each percentile; not in determining skill ranking of different occupations.

Figure 1.2: Polarization Over the Skill Distribution, With and Without Routine Occupations, 1982-2009



Data comes from the March Annual Supplement to the CPS from 1982-2009 for full-time, full-year workers above the age of 16. Figure construction mimics that of Figure 1.1. For the removal of routine occupations, these occupations were removed after the ranking of occupations by mean log wages, but prior to computing growth in wages and employment shares.

removed, a substantial fraction of polarization vanishes from these figures. For employment shares, between the 15th and the 60th percentiles, employment share growth increases substantially when routine occupations are removed. Moreover, with routine occupations removed, employment share growth is positive at all levels of the skill distribution. For wages, the contribution of routine occupations is even more stark. With routine occupations removed, the polarization phenomenon in wages is completely gone. Removing routine occupations leads to wage increases that are monotone in skill.[8] This exercise highlights that recent changes in employment and wage growth in routine occupations contribute substantially to the broader phenomenon of labor market polarization.

[8]Some may find it noteworthy that wage growth actually slight declines for the lowest skill percentile when routine occupations are removed from the sample. As observed in Autor and Dorn (2013), the job requirements of some lower skilled routine occupations (such as some types of secretarial work) may have changed over time and led to wage increases for this small subset of routine occupations, and this would be consistent with these observed patterns.

1.2.1.2 Employment Participation and Skill Levels

It is important to observe, however, that the above on labor market polarization is conditional on employment. But what about labor market outcomes without conditioning on employment, in particular, employment participation decisions? Here I briefly review long run changes in the employment to population ratio across skill populations and show that employment rates, in contrast, have not polarized. As occupation-based measures of skill are not available for non-employed individuals, I measure skill using educational categories.

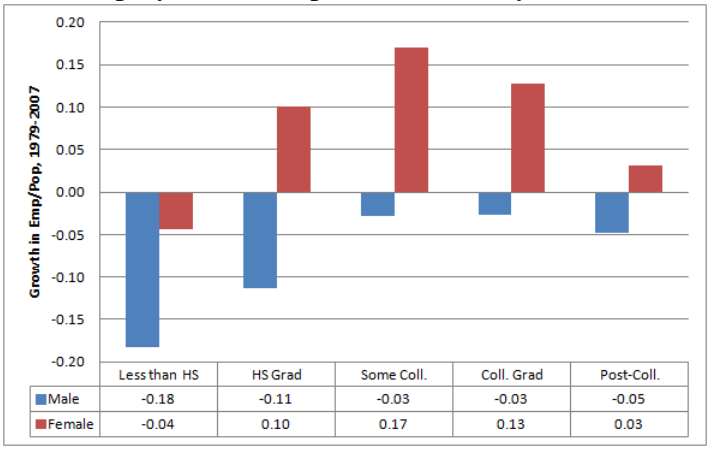
Figure 1.3 uses data from Autor (2010) and presents the growth in the employment to population ratio between 1979 and 2007 for men and women across five different educational groups, where the employment to population ratio is with respect to the population in that skill level, not the aggregate population.[9] I consider men and women separately, as there have been substantial changes in female labor force participation in the 1980s and 1990s. Over this time horizon, the employment to population ratio for men has fallen while the employment to population ratio has risen for women, but both genders show similar patterns in the relative growth across educational groups. The least growth has been observed amongst those with less than a high school diploma, and growth generally increases with educational attainment. For men, this relationship is nearly monotone, with growth in the employment rate increasing with education levels up through post-college education. For women, the relationships is slightly more nuanced. Growth in employment rates increases with education from high school dropouts though having some college, but declines slightly thereafter, particularly for women with more than a college degree. Thus, if anything, for women, middle-skilled individuals have seen the most growth in employment rates.[10]

The evidence of Figure 1.3 stands in stark contrast to the polarization phenomena described before. Although employment growth has declined amongst middle skilled occupations, employment

[9]These measures from Autor (2010) measure employment as being employed at any time in the previous year, and are thus more long-term measures of non-participation.

[10]I give slightly more discussion of this in the following section, where it is evident this is largely due to smaller growth in this ratio during the 1980s and 1990s, when there was rapid expansion of the female labor force participation rate, which may have occurred for many reasons that would interact heterogeneously with skill levels.

Figure 1.3: Growth in Employment to Population Ratio by Educational Level, 1979-2007



	Less than HS	HS Grad	Some Coll.	Coll. Grad	Post-Coll.
Male	-0.18	-0.11	-0.03	-0.03	-0.05
Female	-0.04	0.10	0.17	0.13	0.03

Data is from Autor (2010), drawing from the May ORG Samples of the Current Population Survey. The time window of 1979 to 2007 is chosen such that the starting and endpoint are both at similar points of the business cycle.

rates have fallen the most among the lowest skilled workers, and growth in these rates is generally increasing with the skill level. Thus, explaining the recent changes in inequality across skill levels extends beyond labor market polarization and includes reconciling this disconnect between outcomes conditional and not conditional on employment.

1.2.2 Labor Market Trend Dynamics

Given that background, I now present the four sets of facts that I will use in my quantitative evaluation of the role of routine-biased technological change in labor market outcomes: (1) occupational employment shares (2) relative hourly earnings across occupations (3) employment to population ratios, aggregate and by skill (4) labor's share of income. As understanding the changes occurring in routine occupations is central to polarization, I focus on specific changes in these occupations. From here forward, I restrict my attention to the three occupational groups - abstract, routine, and manual - and I present secular trends in their relative hourly earnings levels and shares of total employment. For all these trends, I focus particularly on the time paths of these changes; this contrasts with much of the existing literature on polarization, which has only focused on changes

between two points in time.[11]

1.2.2.1 Occupational Employment Shares

Figure 1.4 plots the shares of total employment belonging to abstract, routine and manual task occupations from 1976 to 2013. Prior to the 1980s, employment in routine task occupations comprised more than 40% of total employment, the largest employment share of all three occupational groups. But beginning with the early 1980s, employment in routine task occupations has been steadily declining, and in 2013 had fallen to 26% of total employment.[12] Simultaneously, the share of employment in abstract jobs has steadily risen from roughly 25% before the 1980s to well above 40% in 2013. The share of employment in manual jobs has remained roughly stable, hovering at about 30% of total employment throughout the entire sample period, with a slight surge observed in the decade of the 2000s[13].

1.2.2.2 Relative Hourly Earnings Across Occupations

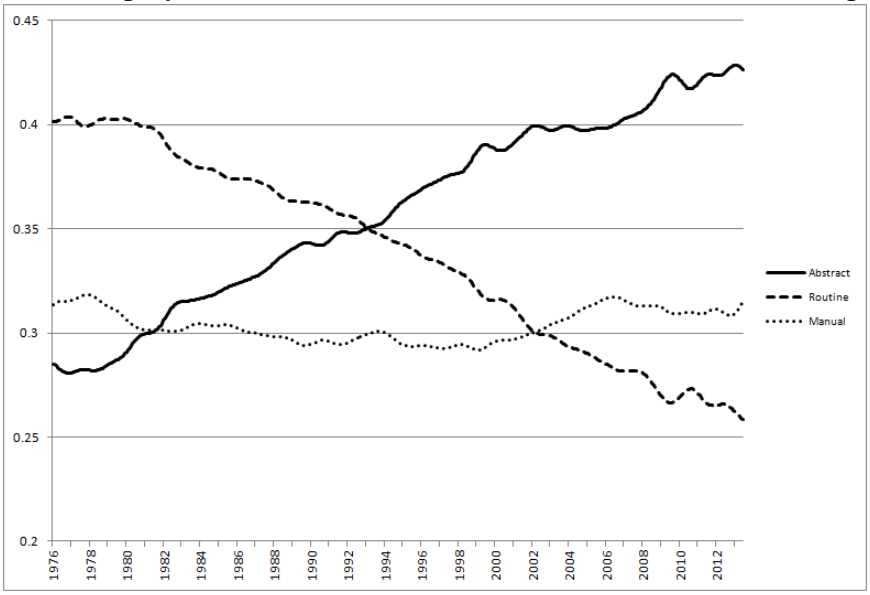
Figure 1.5 presents the average hourly earnings in routine occupations relative to hourly earnings in abstract and manual occupations from 1975 to 2011. Between 1975 and 2011, hourly earnings in routine occupations relative to abstract ones has fallen by 20%. This decline began in the 1980s, but has been particularly substantial in the 1990s and early 2000s, though it appears to have slowed down some in the latter half of the past decade. Meanwhile, hourly earnings in routine occupations relative to manual occupations has has been relatively flat throughout the entire period, with the exception of some slight prior to 1990. These trends are consistent with the evidence already

[11] More recently, Foote and Ryan (2012), Cortes et al. (2012) and Smith (2013) have also given some attention to these time paths.

[12] One common concern is that this decline is simply capturing changes occurring with the manufacturing industry's recent decline. I show in the Appendix that the changes in employment and hourly earnings in routine work are occurring across all industries and that only a third of routine employment is concentrated in manufacturing.

[13] Note that this surge is unlikely to be completely explained by a boom in construction preceding the Great Recession, as much of the growth in the manual employment share has remained since the recession.

Figure 1.4: Employment Shares of Abstract, Routine, and Manual Occupations



Data is from the CPS basic monthly files on employed persons in the non-farm business sector between the ages of 16 and 65. The above results are robust, however, to including the employed over age 65 or including employees in agricultural or public sectors of the economy. Occupational classifications come from Autor and Dorn (2013). Additional details regarding data construction and sources are available in the Appendix.

seen on polarization - not only has employment in routine occupations declined relative to other occupations, but so has hourly earnings.

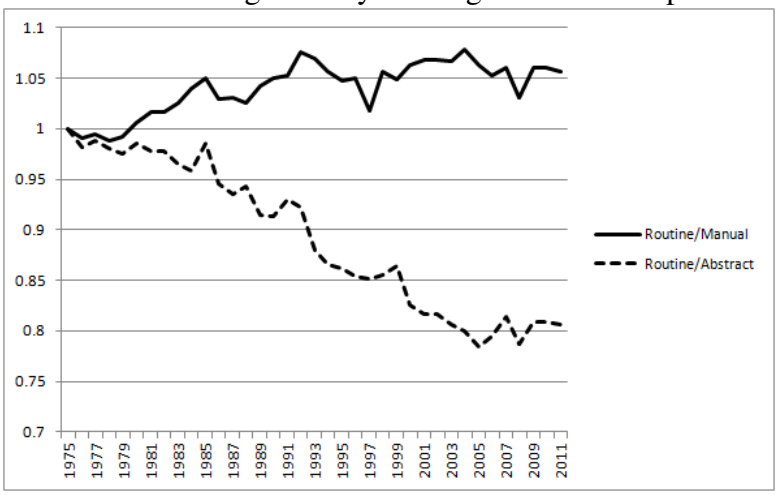
1.2.2.3 Employment to Population Ratio

Figure 1.6 plots the employment to population ratio from 1975 Q1 to 2013 Q3. Through the year 2000, especially in the 1980s, there was an upward trend in the percentage of the population that was employed, largely due to the rise of female labor force participation.[14] However, since the year 2000, there have been substantial declines in this ratio. An ongoing discussion regarding this decline considers whether it is purely due to cyclical factors surrounding the two most recent recessions, or if there are deeper structural forces at work as well.[15]

[14] Although it is true that the participation rate for men has declined over time as well, the decline is quantitatively far smaller (~2%) in the 1980s than is the increase amongst women (~10%).

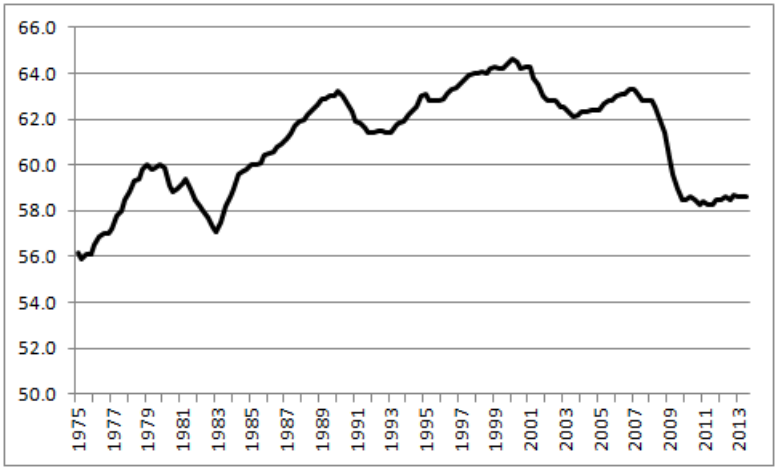
[15] See Jaimovich and Siu (2012) for a recent argument about how routine work could suggest trend movements are embodied in these cyclical fluctuations.

Figure 1.5: Relative Average Hourly Earnings Across Occupational Groups



Data is from the CPS March Supplement files on employed persons in the non-farm business sector. Occupational classifications come from Autor and Dorn (2013). Average hourly earnings is constructed by dividing total wage and salary income for an occupational group by total hours worked by individuals in that occupation. Relative hourly earnings is normalized to 1 in 1975. Additional details regarding data construction and sources are available in the Appendix.

Figure 1.6: Employment to Population Ratio, 1975-2013



Data for the civilian employment to population ratio comes from the FRED database.

As shown above, these changes in the employment to population ratio have not been uniform across skill groups. Figures 1.7 and 1.8 show growth in the employment to population ratio across educational groups for men and women, respectively, for each of the past three decades, ending in 2007 to be distinct from the recent Great Recession.[16] For men, the relationship between growth in the employment to population and skill levels has been fairly consistent throughout time - lower skill levels have seen the most declines throughout each of the past few decades and the amount of decline tends to decrease as skill increases. For women, the picture is slightly less clear. The employment to population rate was growing for almost all skill levels in the 1980s and 1990s, the period when there was a significant expansion of female labor force participation. Fluctuations in the amount of growth across skill levels in this period may be due to differences in the forces generating female labor force participation, be it changes in the welfare program and the earned income tax credit (Meyer and Rosenbaum (2001)) or declining child care costs (Attanasio et al. (2008)).[17] But in the last decade, the relationship between employment rates and skill levels for women has far more resembled patterns observed among men, with greater declines among lower skill levels.

This evidence suggests that, with the exception of changes in female labor supply in the 1980s and 1990s, declines in the employment rates in the past three decades have consistently been the most severe among lower skilled workers.

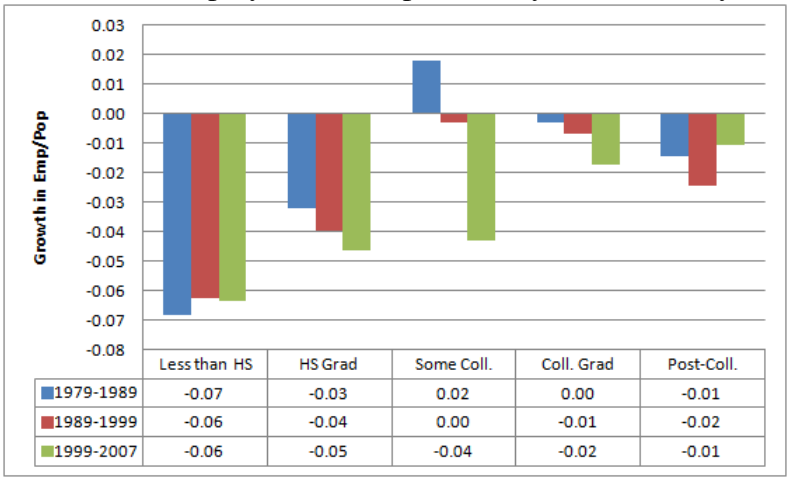
1.2.2.4 Labor Share of Income

Since Kaldor (1961), one of the standard growth facts is that labor's share of total income is constant over time. While this empirical regularity held for many decades in the United States,

[16] Although, as the recent recession has been particularly severe for low-skilled individuals, inclusion of this data would only heighten the results presented here.

[17] It may also simply be that the employment rate for the highest educated women was so high that there was little room for increase. See Juhn and Potter (2006) or Fogli and Veldkamp (2011) for further discussion on the range of possible theories for this increase and how it might differentially affect women of varying skill levels.

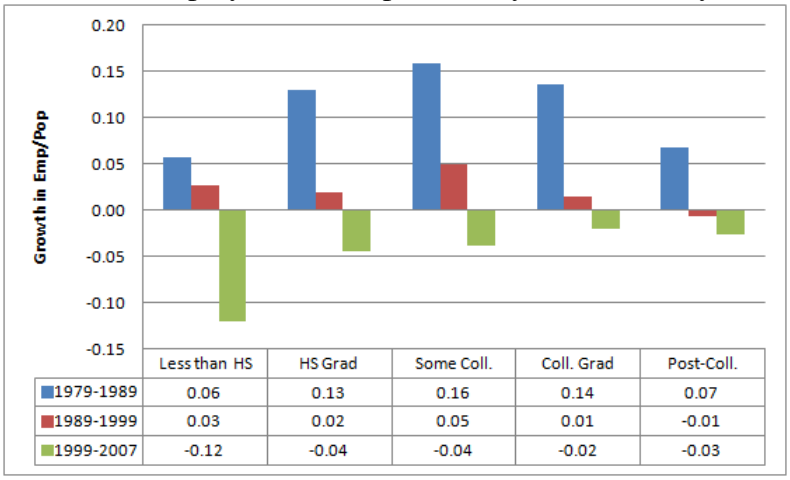
Figure 1.7: Growth in Employment to Population by Skill Level by Decade, Men



	Less than HS	HS Grad	Some Coll.	Coll. Grad	Post-Coll.
1979-1989	-0.07	-0.03	0.02	0.00	-0.01
1989-1999	-0.06	-0.04	0.00	-0.01	-0.02
1999-2007	-0.06	-0.05	-0.04	-0.02	-0.01

Data is from Autor (2010), drawing from the May ORG Samples of the Current Population Survey.

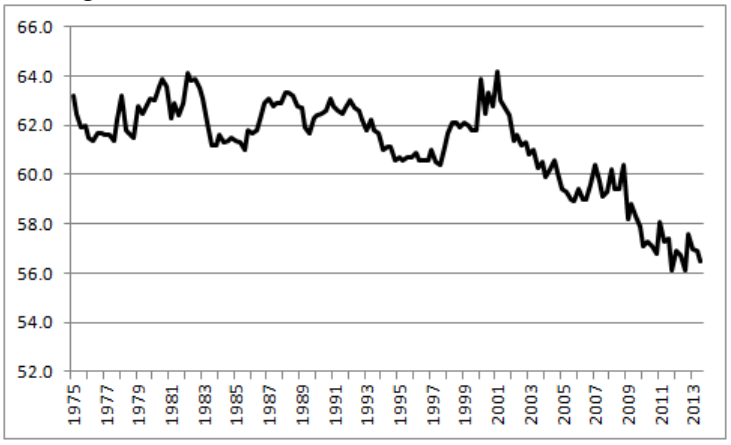
Figure 1.8: Growth in Employment to Population by Skill Level by Decade, Women



	Less than HS	HS Grad	Some Coll.	Coll. Grad	Post-Coll.
1979-1989	0.06	0.13	0.16	0.14	0.07
1989-1999	0.03	0.02	0.05	0.01	-0.01
1999-2007	-0.12	-0.04	-0.04	-0.02	-0.03

Data is from Autor (2010), drawing from the May ORG Samples of the Current Population Survey.

Figure 1.9: Labor's Share of Income, 1975-2013



Data for the labor share of income for the non-farm business sector comes from the BLS Labor Productivity and Costs series.

recent evidence suggests that the labor share has actually been declining.[18] Figure 1.9 presents the labor share of income for the non-farm business sector from 1975 Q1 to 2013 Q3. In the past 40 years, the labor share has fallen by roughly 7 percentage points, about a 10% decline in total.[19] Of note, this decline appears to have accelerated recently, with a much steeper decline since the early 2000s, and as of 2013, there are no signs of a reversal in this trend.

With the addition of this evidence on the labor share, the evidence presented regarding the decline of routine work can be used to disaggregate the labor share into occupational labor income shares.[20] While this does not provide any new evidence per se, it provides a direct link between declines in the labor share and the decline of routine work. Figure 1.10 plots the occupational labor shares - total income paid to an occupation divided by nominal GDP - for abstract, routine and manual occupations. What is immediately evident is that while the aggregate labor been declining, particularly in the 2000s, the shares of income paid to different occupations have been changing quite a bit over time.[21]

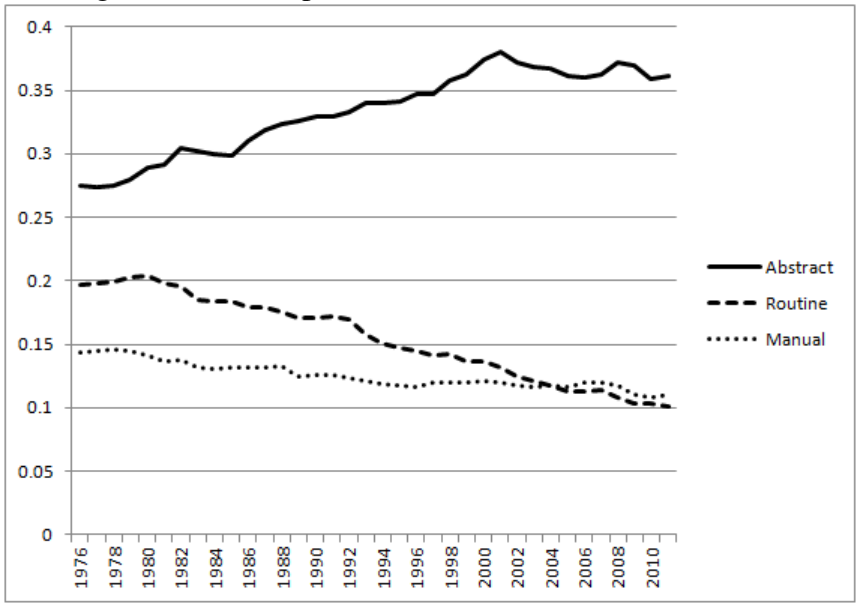
[18] As documented in Karabarbounis and Neiman (2014), this decline is occurring worldwide as well.

[19] Elsby et al. (2013) suggest that a modest fraction of this decline may be due to measurement methods employed by the BLS.

[20] Precise details on how occupational labor income shares are generated are available in the Appendix.

[21] An interesting observation is that most stark change among labor shares in the 2000s, when

Figure 1.10: Occupational Shares of Income, 1976-2011



Occupational income data comes from the March CPS Supplement and data on nominal GDP comes from the BLS Labor Productivity and Costs series. Data is for the non-farm business sector. The occupational share of income for occupation j is defined as total wage and salary income paid to a worker in occupation j divided by nominal GDP. As such, the sum of all occupational shares of income equals the aggregate labor share of income. Additional details regarding data sources and construction are available in the Appendix.

1.3 Model

With the facts I am considering for the quantitative evaluation documented, I now turn to developing the model used for this exercise. In this section, I present an extension of the standard one sector growth model that includes the following three key features: (i) the tasks framework of Autor and Dorn (2013), (ii) a Roy model of occupational choice, and (iii) equipment specific technological change as in Greenwood et al. (1997). Adding these three features generates a rich and flexible framework in which to consider the changes documented in the prior section. Time in the model is discrete, and one period is assumed to equal one year.

1.3.1 Environment

1.3.1.1 Technology

An aggregate production function produces output using five inputs - structures (K_{st}), equipment (K_{et}), and efficiency units of manual (N_{mt}), routine (N_{rt}), and abstract (N_{at}) labor:

$$Y_t = K_{st}^{\alpha}\left[\mu_m N_{mt}^{\sigma} + (1-\mu_m)\left[(1-\mu_a)\left[(1-\mu_r)K_{et}^{\gamma} + \mu_r N_{rt}^{\gamma}\right]^{\frac{\rho}{\gamma}} + \mu_a N_{at}^{\rho}\right]^{\frac{\sigma}{\rho}}\right]^{\frac{(1-\alpha)}{\sigma}} \tag{1.1}$$

The terms in square brackets represent three levels of aggregation. The innermost square brackets combine routine labor and equipment, to produce routine tasks, with the elasticity of substitution between routine labor and equipment given by $\frac{1}{1-\gamma}$. Routine tasks are then combined with abstract labor, with an elasticity of substitution given by $\frac{1}{1-\rho}$. Finally, the combination of routine and abstract tasks is combined with manual labor to complete the CES structure of production, with the elasticity of substitution between manual labor and the routine-abstract composite given by $\frac{1}{1-\sigma}$.

This production technology is a natural generalization of the tasks framework exposited in

the aggregate labor share shows the most decline, was the slowdown of the labor share for abstract occupations. This is a potentially useful insight in understanding the recent labor share decline, and is addressed in Chapter 2.

Autor and Dorn (2013).[22] They argue that if capital equipment and routine labor are substitutes (i.e. $\gamma > 0$), routine tasks and abstract labor are combined in a Cobb-Douglas aggregate (i.e. $\rho = 1$), and manual labor is complementary to the combination of routine and abstracts tasks (i.e. $\sigma < 0$), then this framework provides qualitative predictions consistent with the decline of routine work.[23] Relative to Autor and Dorn, I add structures as an input to production and do not impose $\rho = 1$.

This production structure is also similar to the capital-skill complementarity structure used in Krusell et al. (2000). The capital-skill complementarity considers only two skill groups, skilled and unskilled, however, and is thus not well-suited to a discussion of the decline of routine work, as there is no substitutability between equipment and middle-skilled labor. However, if manual labor tasks are only performed by unskilled labor and abstract labor tasks are only performed by skilled labor, then their framework is the special case of the above where $\mu_r = 0$.

Output can be used for consumption or investment in structures or equipment according to the aggregate resource constraint:

$$Y_t = C_t + I_{st} + I_{et}$$

Following Greenwood et al. (1997), I assume capital accumulates as follows, where q_t represents equipment investment specific technological change:

$$K_{st+1} = (1 - \delta_s)K_{st} + I_{st} \tag{1.2}$$

$$K_{et+1} = (1 - \delta_e)K_{et} + q_t I_{et} \tag{1.3}$$

[22] Which itself builds on earlier work by Autor et al. (2003) and Autor et al. (2006).

[23] Autor and Dorn (2013) consider two output sectors, goods and services, where services is only produced by manual labor and goods are produced by a combination of routine and abstract task inputs. Because household utility is a CES aggregate of both outputs, the final consumption good in their model is the same as mine. Note, however, that their model considers a specialized version of manual tasks where these are exclusively low-skill service occupations, and not manual tasks more generally.

I assume that movement in q_t is exogenous and deterministic. With these laws of motion for capital accumulation, the aggregate resource constraint can be rewritten as:

$$Y_t = C_t + K_{st+1} - (1-\delta_s)K_{st} + \frac{1}{q_t}(K_{et+1} - (1-\delta_e)K_{et}) \tag{1.4}$$

1.3.1.2 Household

There is a representative household comprised of a continuum of individuals.[24] Each individual in the household has preferences over consumption and labor supply, with period flow utility for an agent i given by:

$$U(C_{it}, E_{it}) = log(C_{it}) - E_{it}\chi_{it}^{\nu}$$

where C_{it} represents consumption by agent i, E_{it} is a 0/1 choice regarding work, and χ_i^{ν} is the disutility incurred by working.

Individuals differ in their disutility of work parameter, χ, as well as their skill level, z. Skill is distributed continuously across the real line, with $z \sim F(z)$, and χ follows a uniform distribution, $\chi \sim UNIF(0, \bar{\chi}_t)$. The upper bound for this uniform distribution, $\bar{\chi}_t$, will be allowed to vary exogenously and deterministically over time to proxy for female labor force participation changes driving early growth in aggregate employment rates in the 1980s and 1990s. The skill distribution is assumed to be constant throughout time. I assume that skill level and disutility of work are orthogonal at the individual level.

The household makes decisions regarding who works and who stays at home, and conditional on working, which occupation individuals are employed in. Workers can be employed in one of three types of occupations - abstract, routine or manual. Conditional on employment in occupation

[24] I assume a representative household construct largely for tractability here; a possible extension would be to consider allowing for incomplete consumption insurance across skill types and realistically modeling transfers between agents.

j, $j \in \{a, r, m\}$, a worker with skill level z provides $\phi_j(z)$ efficiency units in production, where $\phi_j(z)$ is an occupation specific productivity function. In equilibrium, earnings of a worker employed in occupation j are given by the product of the efficiency units and the wage per efficiency unit in that occupation: $w_{jt}\phi_j(z)$. I assume that occupational productivity is non-negative for all values of skill in each occupation.

As in Jung and Mercenier (2010) and Cortes (2012), I assume that workers with higher skill have a comparative advantage in abstract jobs relative to routine jobs and in routine jobs relative to manual jobs. Specifically, this is imposed by the following condition, akin to log supermodularity:

$$0 \leq \frac{dln(\phi_m(z))}{dz} < \frac{dln(\phi_r(z))}{dz} < \frac{dln(\phi_a(z))}{dz} \tag{1.5}$$

For notational convenience in presenting the household's maximization problem, I first describe several properties of the optimal decision rules for occupational choice and employment participation. A direct implication of the comparative advantage condition in (1.5) is that it is optimal for the household to choose two skill cutoffs, z_{0t} and z_{1t}, where employed individuals with skill below z_{0t} are employed in the manual occupation, workers with skill between z_{0t} and z_{1t} are employed in the routine occupation, and workers with skill above z_{1t} are employed in the abstract occupation.[25] This implies that the following indifference relationships must hold, characterizing the relationship between relative wages and the occupational choice decision:

$$w_{rt}/w_{at} = \frac{\phi_a(z_{1t})}{\phi_r(z_{1t})} \tag{1.6}$$

$$w_{rt}/w_{mt} = \frac{\phi_m(z_{0t})}{\phi_r(z_{0t})} \tag{1.7}$$

The household will find it optimal to choose a cutoff function, $\hat{\chi}_t(z)$, to determine which workers are employed and which stay at home, where a worker of skill level z is employed only if his

[25] These sorting conditions will have strong implications regarding which types of workers transition in and out of these occupations; Cortes (2012) has shown empirical support for these implications.

disutility of work is less than the cutoff, $\chi < \hat{\chi}_t(z)$. Given that skill and the disutility parameter are orthogonal, similar to Gali (2011), the employment rate for skill level z can be expressed as $\hat{\chi}_t(z)/\bar{\chi}_t$.

The household maximizes an equal weighted integral of individual utilities. Given separability of preferences over consumption and labor supply at the individual level, the household will equalize consumption across individuals and flow utility for the household in period t is given by:

$$\begin{aligned} U(C_t, \hat{\chi}_t(z)) &= log(C_t) + \int_{-\infty}^{\infty} \int_{0}^{\hat{\chi}_t(z)} \frac{\chi^{\nu}}{\bar{\chi}_t} d\chi dF(z) \\ &= log(C_t) + \int_{-\infty}^{\infty} \frac{(\hat{\chi}_t(z))^{1+\nu}}{1+\nu} \frac{1}{\bar{\chi}_t} dF(z) \end{aligned}$$

As shown in Gali (2011), ν is equivalent to the inverse Frisch elasticity of labor supply.

In equilibrium, the household takes prices as given and chooses sequences for consumption, next period's capital stocks, and cutoffs for occupational choice and employment to maximize:

$$\sum_{t=0}^{\infty} \beta^t \left\{ log(C_t) - \int_{-\infty}^{\infty} \int_{0}^{\hat{\chi}_t(z)} \frac{\chi^{\nu}}{\bar{\chi}_t} d\chi dF(z) \right\}$$

subject to their budget equation for each period:

$$\begin{aligned} C_t + K_{st+1} + (1-\delta_s)K_{st} + \frac{1}{q_t}(K_{et+1} - (1-\delta_e)K_{et}) = \\ w_{mt} \int_{-\infty}^{z_{0t}} \phi_m(z) \frac{\hat{\chi}_t(z)}{\bar{\chi}_t} dF(z) + w_{rt} \int_{z_{0t}}^{z_{1t}} \phi_r(z) \frac{\hat{\chi}_t(z)}{\bar{\chi}_t} dF(z) + w_{at} \int_{z_{1t}}^{\infty} \phi_a(z) \frac{\hat{\chi}_t(z)}{\bar{\chi}_t} dF(z) + \\ r_{et} K_{et} + r_{st} K_{st} \end{aligned}$$

as well as subject to non-negativity constraints on capital and consumption and subject to initial capital stocks.

1.3.1.3 Firm

A representative firm produces aggregate output and behaves competitively, renting equipment and structures from the household at rental rates, r_{et} and r_{st}, and paying wages w_{jt} per efficiency units of labor provided in occupation j, $j \in \{a, r, m\}$. The firm's problem is standard and is given by:

$$max_{K_{st},K_{et},N_{mt},N_{rt},N_{at}} \quad \pi_t = Y_t - r_{st}K_{st} - r_{et}K_{et} - w_{mt}N_{mt} - w_{rt}N_{rt} - w_{at}N_{at}$$

1.3.2 Equilibrium

I focus on competitive equilibrium outcomes. A competitive equilibrium is an set of allocations $(C_t, K_{et}, K_{st}, N_{at}, N_{rt}, N_{mt}, z_{0t}, z_{1t}, \hat{\chi}_t(z))$, such that given prices $(r_{et}, r_{st}, w_{at}, w_{rt}, w_{mt})$:

- $(C_t, K_{et}, K_{st}, z_{0t}, z_{1t}, \hat{\chi}_t(z))$ solves the household problem, taking $(w_{mt}, w_{rt}, w_{at}, r_{st}, r_{et})$ and the initial conditions as given
- $(K_{et}, K_{st}, N_{mt}, N_{rt}, N_{at})$ solve the firm problem, taking $(w_{mt}, w_{rt}, w_{at}, r_{st}, r_{et})$ as given
- Labor, Capital and Output Markets clear

In equilibrium, the share of employment and average hourly earnings in occupation j are given by:

$$\frac{E_{jt}}{E_t} = \frac{\int_{Z_{jt}} (\hat{\chi}_t(z)/\bar{\chi}_t) dF(z)}{\int_{-\infty}^{\infty} (\hat{\chi}_t(z)/\bar{\chi}_t) dF(z)}$$

$$w_{jt}N_{jt}/E_{jt} = \frac{w_{jt} \int_{Z_{jt}} \phi_j(z) (\hat{\chi}_t(z)/\bar{\chi}_t) dF(z)}{\int_{Z_{jt}} (\hat{\chi}_t(z)/\bar{\chi}_t) dF(z)}$$

where $Z_{jt} = \{z \mid \text{skill } z \implies \text{employment in occupation } j \text{at time } t\}$. It is clear that average hourly earnings in a given occupation is determined by two channels: the level of the wage and the set of skills in that occupation. I will revisit these two margins when I present the analysis of the quantitative results.

1.4 Driving Forces and Solution Method

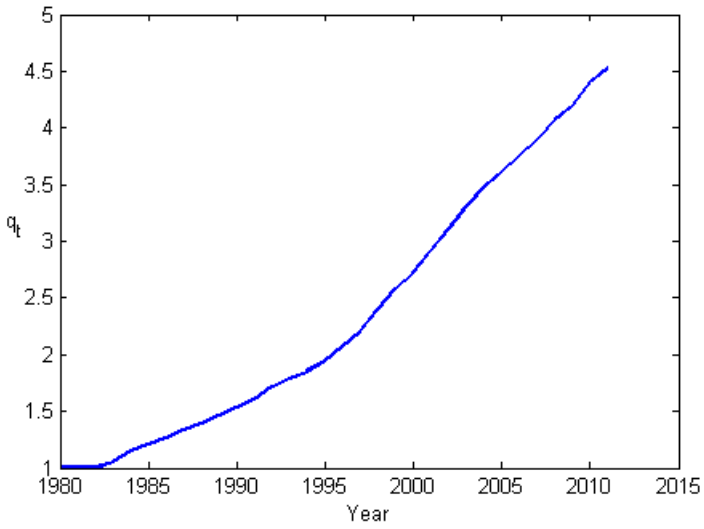
The model described in the previous section features two exogenous driving forces: q_t and $\bar{\chi}_t$. In this section, I describe how these will be measured from the data and how I solve for the competitive equilibrium.

1.4.1 Measuring Driving Forces

The equilibrium resource constraint in (1.4) implies that equipment investment specific technological change, q_t, is equivalently interpreted the amount of investment that can be purchased with one unit of consumption - the relative price of equipment investment. Thus, as in Greenwood et al. (1997) and Krusell et al. (2000), I measure q_t as the price of consumption divided by the price of equipment investment. To measure the price of equipment investment, I utilize the equipment price series of DiCecio (2009), who extends through the year 2011 the quality bias price corrections to the BEA's investment price index performed originally by Cummins and Violante (2002) and Gordon (1990). To measure the price of consumption, I use BEA data on the price indices for non-durable consumption and non-housing and non-energy services, and then apply quality bias adjustments to these series as suggested in Boskin et al. (1993) and Gordon (2006). Details of these bias adjustments are provided in the Appendix. The ratio of the consumption price divided by the equipment price yields a series for q_t through the year 2011. Figure 1.11 shows q_t from 1980 to 2011.

A relevant question here is whether using the price of total equipment is the most appropriate measure in computing technological change. Most theories of technological change and routine work postulate that specifically ICT capital substitutes for routine labor, as compared to all equipment in general. Not wanting to take a strong stance on which specific types of equipment are more or less substitutable for routine labor, I simply use total equipment as a benchmark for this analysis. That said, as mentioned in Cummins and Violante (2002), growth in the technological change measure, q_t, from the early 1980s onward is substantially driven by both the increased share of investment in and the price decline of information and communications technology (ICT)

Figure 1.11: Time Path for q_t



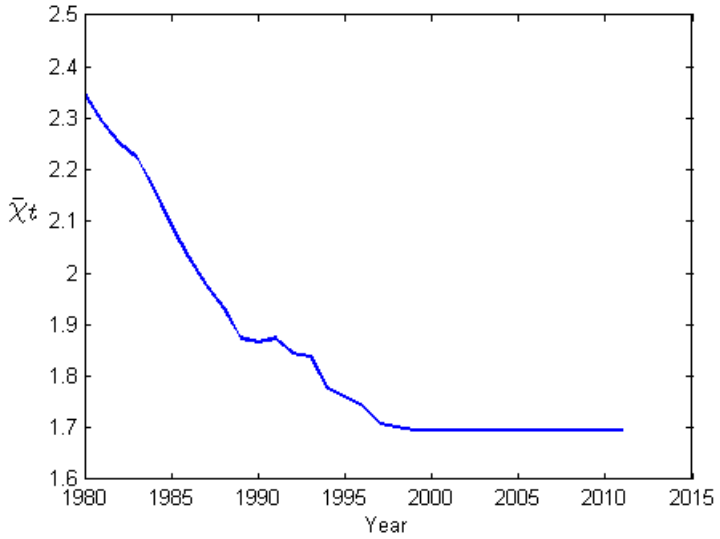
The series for q_t is generated using quality adjusted equipment prices from DiCecio (2009) and quality adjusted consumption prices for non-housing, non-energy services and non-durable consumption. Unadjusted consumption prices are available from the BEA; the adjustment for quality bias is detailed in the Appendix.

equipment, suggesting that this benchmark may not be so bad relative to a narrower definition of equipment.[26]

As changes in $\bar{\chi}_t$ are used to proxy for factors leading to increased female labor force participation, I choose the time path of $\bar{\chi}_t$ to mimic the time path of female participation rate. In practice, given an initial and terminal value for $\bar{\chi}_t$, I scale the inverse of growth in female labor force participation in each period to match the total change in $\bar{\chi}_t$ - additional details of how this is performed are available in the Appendix. Growth in $\bar{\chi}_t$ stops in the year 1999, when female labor force participation attains its maximum in the postwar period. The initial value of $\bar{\chi}_t$ is calibrated so that the employment to population ratio in at the model's initial point matches the data in 1980 and the terminal value of $\bar{\chi}_t$ is calibrated along the transition path to match changes in the aggregate employment rate over the same time period used for calibrating the elasticity parameters. The time

[26] Although there is also growth in the q_t series in the years prior to the 1980s, because this growth is driven by other types of equipment than ICT technology, I consider q_t in these prior years to represent a very different measure of technological change. As such, I do not use evidence on q_t for these prior years in any of the empirical implementation.

Figure 1.12: Time Path for $\bar{\chi}_t$



The time path for $\bar{\chi}_t$ is generated using the inverse of the female labor force participation rate between 1980 and 1999. The precise method for fitting the female labor force participation rate to the data is described in the Appendix.

path of $\bar{\chi}_t$ is plotted in Figure 1.12.

1.4.2 Solution Method

Although the data series for q_t is only available through the year 2011, the model's results for the year 2011 and earlier will depend on what agents believe the future time path is for q_t, as its future path will influence agents' decisions along throughout the model's time paths. One possible approach is to allow q_t to continue to grow at some constant rate along a balanced growth path, however, as discussed in Greenwood et al. (1997), a balanced growth path will require production to Cobb-Douglas so q_t can be expressed as labor-augmenting technical change.[27] As the facts regarding employment and hourly earnings growth in occupations generate substantial secular trends in factor shares, requiring production to be Cobb-Douglas is unappealing in this

[27]See He and Liu (2008) for an additional discussion on this point. Maliar and Maliar (2011) present a model whose balanced growth path allows for positive growth in q_t and CES production by introducing other sources of growth. However, for the existence of a balanced growth path, these other source of growth must generate constant factor shares, as in the Cobb-Douglas case.

context. Thus, I do not use an assumption of balanced growth to anchor the model's asymptotic behavior.

Instead, I assume that equilibrium objects asymptotically converge to a final steady state where all variables are constant. Namely, I view the changes in q_t and $\bar{\chi}_t$ as inducing a transition of the economy from one steady state to another. I constrain growth in q_t after 2011 to gradually slow until it reaches a constant value after a finite number of periods; in practice, I fit a second order polynomial to the end of the observed series and project forward gradually declining growth through the year 2019, after which q_t is constant. However, the results are robust to slowing growth more gradually and constraining q_t to be constant at a much later date. With q_t reaching a terminal value, the model will converge asymptotically to a new steady state. Numerically, I solve for the transition path from the initial conditions to this final steady state using a multiple shooting algorithm, where I impose that the convergence occurs after a fixed number of periods.[28] Additional details regarding the numerical solution are available in the Appendix. To determine the initial conditions, I start the model in steady state in the year 1980 with q_t normalized to 1 and $\bar{\chi}_t$ at its initially calibrated value.

1.5 Calibration

I now turn my attention to calibration of the model's parameters. Intuitively, the three parameters governing elasticities of substitution between different labor inputs and capital equipment - γ, ρ, and σ - will have crucial implications for the model's predictions regarding the decline of routine work. Autor and Dorn (2013) argued how these elasticities of substitution influence the qualitative changes in employment and hourly earnings across different occupations and much of the intuition they develop carries over to this context.[29] There are a number of delicate issues in calibrating these parameters, so I momentarily defer a discussion of that calibration and first present the calibration

[28] In the results presented, I check for convergence 25 years following the end of growth in q_t in 2019. But the results are robust to checking for convergence at a longer time horizon, such as 75 years or more.

[29] As there are some similarities with the capital-skill complementarity framework, the intuition developed in Krusell et al. (2000) is also applicable in this context.

of the other remaining parameters.

1.5.1 Calibration of Other Parameters

Aside from the elasticity of substitution parameters, I need to calibrate the Cobb-Douglas parameter on structures (α), the share parameters in the production function (μ_j), the depreciation rates for structures and equipment capital (δ_s, δ_e), the discount factor (β), the inverse Frisch elasticity (ν), and the functional forms regarding occupation specific productivities ($\phi_j(z)$) and the skill distribution ($F(z)$), with any attendant parameters. Full details regarding the calibration strategy and measurement methods taken to calibrate these are contained in the Appendix; the following reports the general strategy.

The Cobb-Douglas structures parameter, depreciation rates, discount factor and inverse Frisch elasticity can all be calibrated without numerically solving the model. In the model, the share of income paid to structures is given by the Cobb-Douglas parameter; α is thus set to the average share of income in structures in the data for the sample period of 1980-2011. In the model's steady state, the depreciation rates are equivalent to effective investment rates; δ_s and δ_e are thus set to the average effective investment rates in the data for structures and capital, respectively, over 1980-2011. I choose standard values for the discount factor and the inverse Frisch elasticity, $\beta = 0.96$ and $\nu = 0.5$.

For the skill distribution and occupation specific productivities, I choose tractable functional forms that are also consistent with the assumptions imposed earlier. Occupational productivity functions take an exponential form, $\phi_j(z) = e^{a_j z}$, consistent with the comparative advantage condition in (1.5). I assume that worker skill follows a normal distribution, with $z \sim N(0, \sigma_z^2)$. The choice of a normal distribution seems to be a reasonable benchmark; the implication of this choice is that the log income distribution will be piecewise log normal, which provides a reasonable approximation to the total income distribution.

With these choices for functional forms, I calibrate the share parameters in production (μ_j), occupation productivity parameters (a_j), and the variance of the skill distribution (σ_z^2) by numeri-

cally solving the steady state used to determine the initial conditions, and matching certain model moments in said steady state to those observed in the data in the year 1980.[30] The share parameters in production are calibrated to match the initial levels of occupational shares of income. I calibrate σ_z^2 to match the cross-sectional variance of log earnings for workers in abstract occupations.[31] Calibrating the occupation productivity parameters, however, requires some sensitivity, as they need to satisfy the restriction $a_m < a_r < a_a$ to be consistent with (1.5). Two of the three moments used are the initial employment shares in routine and manual occupations. With these two moments determining two of the productivity parameters, this imposes a fixed set of possible parameter values for all three parameters. In experimenting with different values within this set of parameter, I find that the model's results are robust to picking any of the parameter values in this set. For the results shown below, I choose the set of parameters that calibrates the model to match the initial cross-sectional variance of routine worker log earnings relative to the cross-sectional variance of manual worker log earnings.

1.5.2 Calibrating Elasticity Parameters

For values of the elasticities of substitution between abstract, routine and manual tasks and equipment, ideally there would be existing (preferably consensus) estimates in the literature to draw from. Unfortunately, there have not been been any attempts yet to estimate these specific parameters. In the absence of such values, as an initial benchmark, I pursue a calibration approach where I select values for these three elasticity parameters to match changes in employment shares and growth in the the capital equipment to employment ratio.[32] These changes seem to be rea-

[30] This calibration also requires the calibrated values of the elasticity parameters, discussed in the next section. The full numerical process to simultaneously calibrate all parameters is presented in the Appendix.

[31] Alternatively, I can target the total variance of log income for all workers with no difference on the final results.

[32] An alternate approach would be to estimate these parameters from the production function using data on both prices and quantities. This is slightly more difficult for my model, as efficiency units of labor are not directly observable, and in the model, will be significantly affected by the employment participation decision. As estimation would then require estimating the entire model,

sonable targets, as the elasticity parameters will strongly affect movements in these data moments. As as consequence, with no calibration restrictions placed on prices, the model's predicted outcomes for changes in average hourly earnings over the entire sample can be used in quantitatively evaluating technological change as a driver of the decline of routine work.

In the time period used for calibration, the model will attribute all movements in employment shares and in the ratio of capital equipment to employment to the interaction of technological change (and a falling upper bound for the worker draws determining disutility of work) with the shape of the production function. This naturally biased the model's results towards technological change, and could be problematic if there are other factors outside the model which are influencing these factor shares through a different channel. One such prominent factor that could be influencing employment shares and capital equipment growth is the rise of globalization and the corresponding increased international trade that has occurred over the past several decades.[33]

However, as discussed in Katz and Autor (1999) and Autor, Dorn, and Hanson (2013a), in the decade of the 1980s, trade and globalization factors seemed far less important in shaping labor markets outcomes. Figure 1.13 reinforces this idea by plotting two measures of international trade and globalization, the share of imports in GDP and the ratio of foreign direct investment abroad to domestic non-residential investment, from 1960 to 2011. Although there has been dramatic growth in both of these measures in the past several decades, the decade of the 1980s does stand out as a period of lower levels of and fewer changes in trade and globalization activity. In the left panel of Figure 1.13, we see that although the share of imports in GDP roughly doubled in the 1970s, and subsequently increased by another 60% in the 1990s and 2000s, the share is almost constant throughout the entire 1980s and early 1990s, only jumping upward again in 1994. In the right panel of Figure 1.13, a similar pattern is observed for the ratio of US foreign direct investment abroad to domestic investment. US FDI abroad was 5% that of non-residential domestic investment prior to

I find this approach less desirable. However, I am working on alternate approaches to estimating these parameters.

[33]This channel is highlighted in Autor, Dorn and Hanson (2013a), who show how variation in the exposure to trade shocks from China is correlated with changes in the occupational composition and levels of employment across local labor markets.

the 1980s, and then actually falls in the 1980s, reaching a sample low in 1983, and never recovering to its pre-1980s level until the early 1990s, jumping upward in the year 1993. From there, this ratio has quadrupled to 20% of non-residential domestic investment in 2012, but this growth has taken place in the 1990s and 2000s. The timing of these changes contrasts with the timing of technological change measured by q_t, presented earlier in Figure 1.8, where it was observed that growth in q_t was particularly pronounced both in the 1980s and beyond.

I take advantage of this time window in the 1980s and the early 1990s when trade forces appear to be less prominent in determining labor market outcomes and restrict the calibration of elasticity parameters to only match changes occurring between 1985 and 1992. Specifically, I choose the three elasticity parameters so that the change in the routine employment share, the change in the manual employment share, and the growth in the ratio of equipment to total employment that occurs along the model's transition path between the years of 1985 and 1992 exactly match the observed changes in the data for that same time horizon.[34]

I use 1992 as the endpoint for my calibration window, as the evidence in Figure 1.13 suggests that trade and globalization forces began increasing significantly in the years 1993 and 1994.[35] I use 1985 as the starting point for my calibration window to avoid any irregularities stemming from the assumption that the economy is in steady state in 1980. A consequence of this timing for the calibration procedure is that the model's predicted outcomes regarding employment shares for the period after 1992 can be considered out of sample forecasts for the effects of technological change, and also be used as evidence in the quantitative analysis.

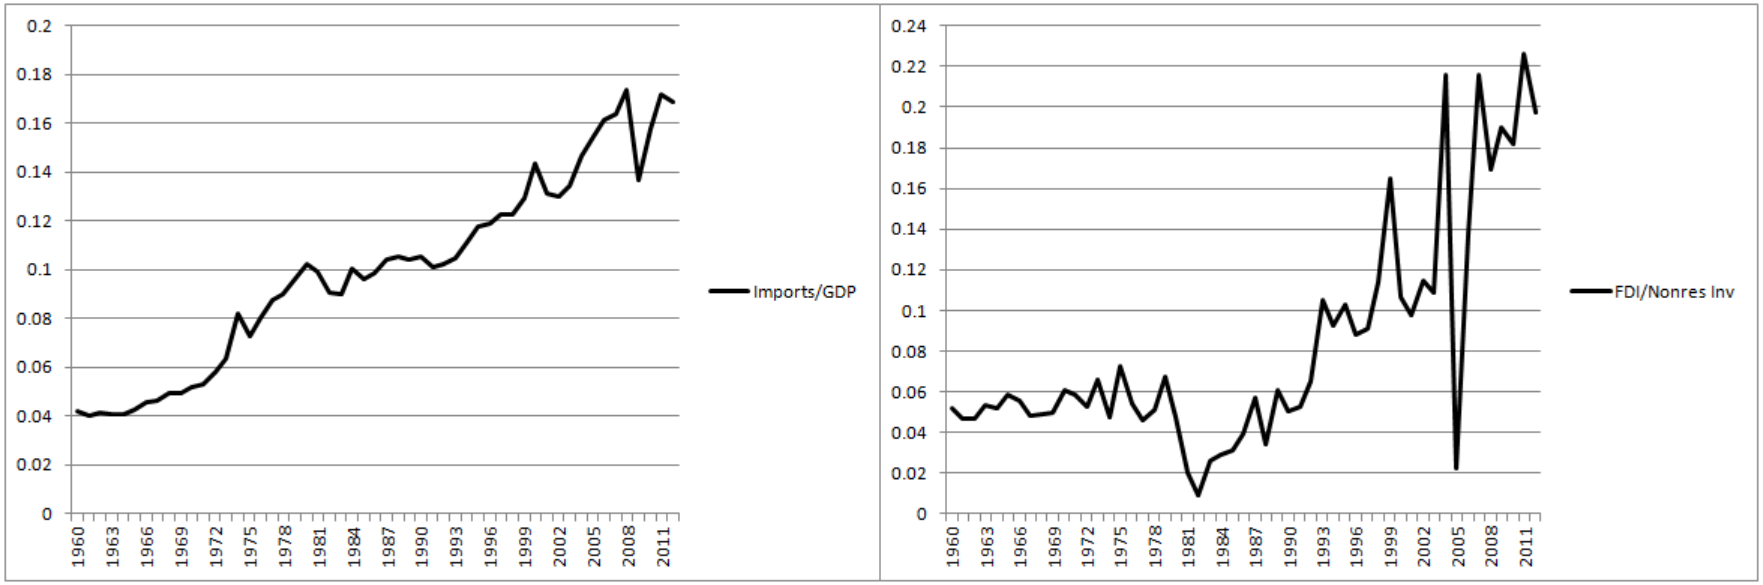
1.5.3 Parameter Values

Table 1.1 lists the calibrated values for all model parameters. Parameters calibrated in steady state to match data moments are able to match the targets to an arbitrary level of precision; details regarding the the exact match are available upon request.

[34] In practice, I match the smoothed trends of these statistics to eliminate business cycle frequency fluctuations in calibrating these parameters. Greater details for this calibration exercise are given in the Appendix.

[35] However, there are not significant differences in ending slightly earlier, in 1990.

Figure 1.13: Measures of Increasing Global Activity: Imports to GDP ratio and US FDI to Domestic Investment Ratio



The left panel presents the ratio of total imports to GDP; the right panel presents the ratio of U.S. private foreign direct investment abroad to U.S. private non-residential investment. Data on imports and GDP comes from Table 1.1.5 of the National Income and Product Accounts at the BEA. Data on private non-residential investment comes from Table 5.3.5 of the National Income and Product Accounts. Data on US Private Foreign Direct Investment Abroad comes from the BEA's U.S. International Transactions data. All ratios are in nominal terms.

Table 1.1: Calibrated Parameters and Targets

Parameter	Value	Target
σ	0.5190	Growth in Capital Equipment per Worker, 1985-1992
ρ	-0.7948	Change in Routine Emp. Share, 1985-1992
γ	0.2728	Change in Manual Emp. Share, 1985-1992
α	0.1651	Average Structures Factor Share
μ_m	0.2564	1980 Manual Labor Factor Share
μ_a	0.5485	1980 Abstract Labor Factor Share
μ_r	0.5418	1980 Routine Labor Factor Share
δ_s	0.0480	Avg. Structures Investment Rate
δ_e	0.1704	Avg. Equipment Investment Rate
a_m	0.0539	Relative Volatility of Log Income - Routine/Manual, 1980
a_r	0.0736	Routine Emp. Share, 1980
a_a	0.7332	Abstract Emp. Share, 1980
ν	0.5	Standard
β	0.96	Standard
σ_z^2	1.1787	Log Income Volatility - Abstract, 1980
$\bar{\chi}_{init}$	2.3466	Employment to Population, 1980
$\bar{\chi}_{term}$	1.6948	Change in Employment to Population, 1985-1992

Details of measurement for all data targets are available in the Appendix.

The values obtained for the three elasticity parameters are $\gamma = 0.2728$, $\rho = -0.7948$, and $\sigma = 0.5190$. Although there are no existing estimates of these elasticities, there are similar elasticities which can be used as a comparison. For example, as mentioned in section 3, my model has some similarities to the capital-skill complementarity framework of Krusell et al. (2000) and if manual and abstract labor is comparable to their unskilled and skilled worker categories, their estimated elasticities between skilled labor and capital equipment and between skilled and unskilled labor may relate to the parameters ρ and σ, respectively. For their parameters similar to ρ and σ, Krusell et al. obtain estimates of -0.495 and 0.401, which are not too far off from the values of -0.7948 and 0.5190 obtained here, and my estimate for σ is well within a standard deviation of the Krusell et al. estimate. Further, these values are within the range of values of elasticities between skilled and unskilled workers and between skilled workers and equipment reported in Hamermesh (1993). Thus, these parameters obtained through the calibration strategy seem reasonable as an initial benchmark.

1.6 Results

Here, I present the model's results opposite the data for the four facts described prior: (1) occupational employment shares (2) relative hourly earnings across occupations (3) employment to population rates (4) labor share of income. Prior to presenting these results, however, I provide some intuition for the results. I then return to some of this intuition in giving some analysis of the results after reporting the model's performance for these four sets of facts.

1.6.1 Initial Intuition

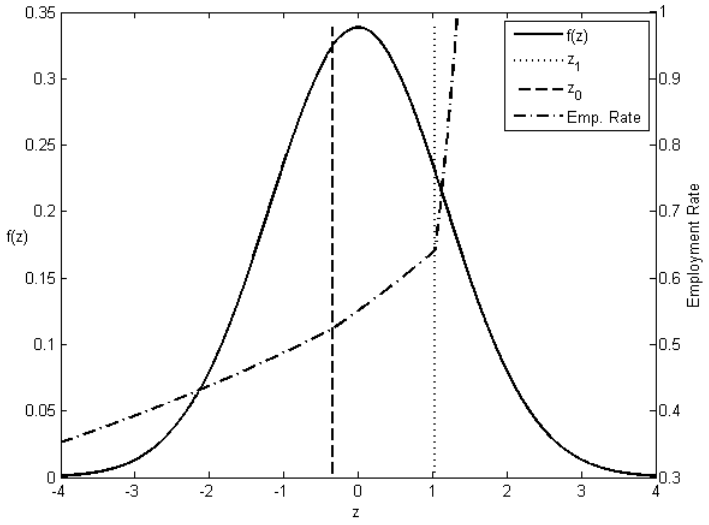
The central mechanism of the model is the production technology, through which changes in the price of equipment capital will affect relative wages and labor demand across different occupations. However, the intuition for how these labor demand effects take place is fairly straightforward and has been discussed substantially in other places, such as Autor and Dorn (2013) and Krusell et al. (2000). The simple idea is that changes in q_t lead to equipment accumulation, which leads to changes in marginal products of production inputs, with the elasticities of substitution shaping how changes in capital accumulation translate to changes in relative labor demands and relative wages. Given the production structure and the calibrated parameter values above, it is clear that changes in q_t will lead to a substitution away from routine labor to equipment, and to increased demand for abstract labor, with ambiguous effects on manual labor. What remains is to quantify these changes in employment and wages.

However, it is helpful to gain some initial insight regarding the household side of the problem. The key margins on the household side are the sorting thresholds for occupational choice (z_{0t} and z_{1t}) and the employment rate ($\hat{\chi}_t(z)/\bar{\chi}_t$). As shown in (1.6) and (1.7), the sorting thresholds in occupational choice are functions of the relative wages, and the household's optimization for the employment rate gives the following expression:

$$\hat{\chi}_{jt}(z)/\bar{\chi}_t = min\left(\frac{1}{\bar{\chi}_t}\left(\frac{w_{jt}\phi_j(z)}{C_t}\right)^{\frac{1}{\nu}}, 1\right) \tag{1.8}$$

Figure 1.14 plots the model results for the employment rate and sorting thresholds against

Figure 1.14: Employment Rate and Sorting Thresholds, 1985



$f(z)$ is plotted against the left y-axis; the employment rate, $\hat{\chi}_t(z)/\bar{\chi}_t$ is plotted against the right y-axis.

the underlying skill distribution in the year 1985, and shows how the continuum of workers are allocated across occupations and employment participation. I will return to this figure following the analysis of the results, to highlight how the results are obtained from movements in these thresholds and the employment rate.

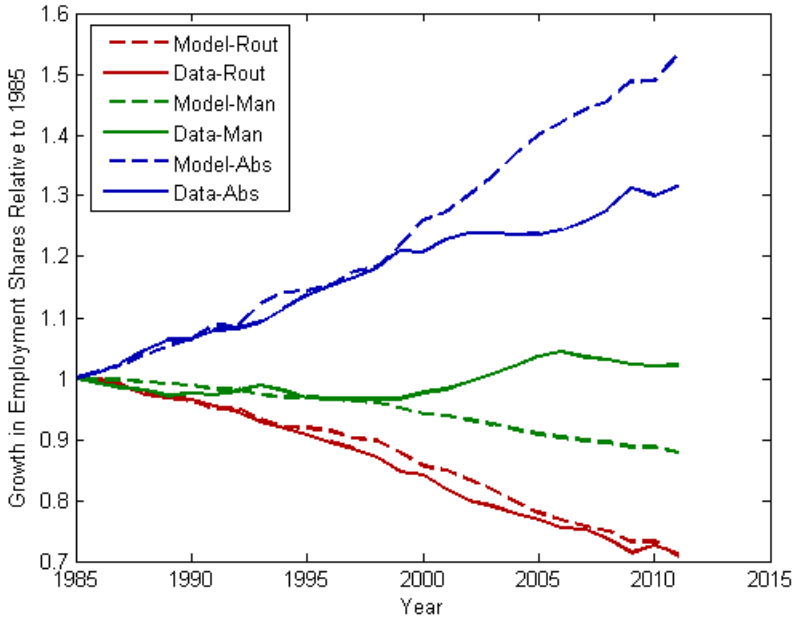
1.6.2 Model Predicted Outcomes

1.6.2.1 Employment Shares

Figure 1.15 presents employment shares in abstract, routine and manual occupations from the data and the model for 1985-2011, where employment shares have been normalized to 1 in the year 1985. Consistent with the calibration, the model matches the trend changes in all three employment shares very well for the years 1985-1992. Further, quantitatively and qualitatively, the model results for the time period of 1992-2000 match almost exactly the observed employment share trends in the data.

From 2000 on, however, the model's predictions are inconsistent with the observed patterns in the data. The model's share of employment in abstract occupations continues to grow and even

Figure 1.15: Employment Shares 1985 to 2011, Model vs. Data



The employment share for routine occupations is in red manual in green, and abstract in blue. Model outcomes are denoted with dashed lines. Annual data on employment shares is obtained by taking 12 month averages of employment shares from the monthly employment shares obtained from the CPS monthly files, whose construction is described in the Appendix. All series are normalized to 1 in 1985.

accelerates in this last decade, while there is an accelerated decline in the model's share of manual workers. This contrasts with the nonlinear patterns observed in the data, as the employment share in abstract occupations flattens out for much of the decade and the share of employment in manual occupations grows steadily. Thus, even with the calibration strategy of matching the employment shares exactly in the 1980s, the model is unable to reconcile these recent changes, as growth has continued (and even accelerated) in q_t in this later period but growth in abstract employment shares has not followed suit.

That said, although the pace of this change is inconsistent with the behavior of abstract and manual employment shares, it is consistent with the continuing decline of the routine share of employment. The model's predictions regarding the routine employment share match the data throughout the 2000s, and entirely explain the total decline in the routine employment share. Thus, the pace of technological change is consistent with the pace of the decline of routine employment; it is simply that the relative growth in other occupations has not followed a similar pattern.

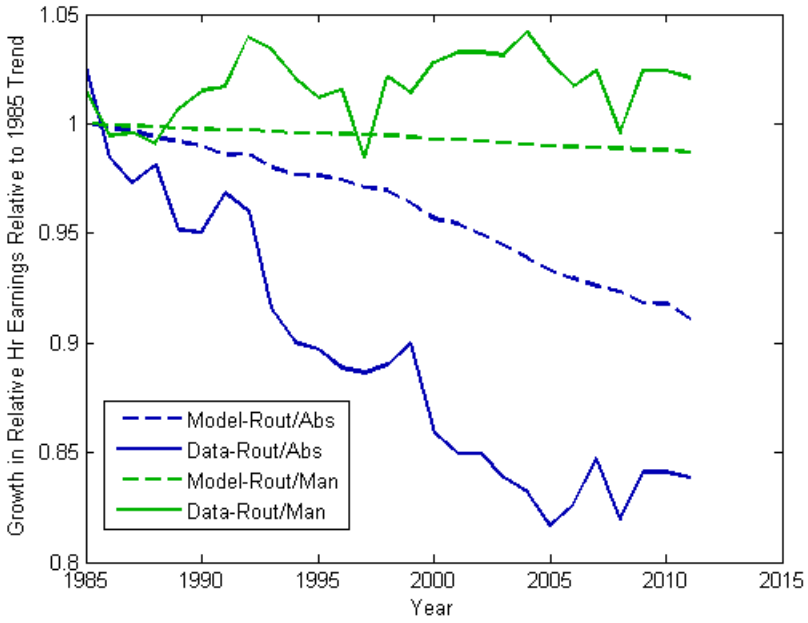
1.6.2.2 Relative Hourly Earnings

Figure 1.16 presents the relative average hourly earnings between occupational groups in the model and in the data. As there are substantial cyclical fluctuations in the data, I normalize the data and the model results to their trend values in 1985.[36]

The model's predictions regarding trends in hourly earnings in routine occupations relative to abstract ones are generally qualitatively consistent with the data, but differ quantitatively. Both the model and the data show a steady decline in routine hourly earnings relative to abstract hourly earnings, with the model explaining roughly 55% of this decline (8.9% out of 16.0%) between 1985 and 2011. However, the relative hourly earnings series in the data actually appears to level out after 2005, whereas the model continues to show a decline, though this is likely to be heavily affected by cyclical factors connected with the Great Recession. If one compares the decline in relative hourly earnings between 1985 and 2000 or 2005, the model is only able to capture 35-40%

[36]I compute the trend for both series using the Hodrick Prescott filter with a parameter of 6.25.

Figure 1.16: Relative Hourly Earnings Between Occupations 1985 to 2011, Model vs. Data



Average hourly earnings in routine occupations divided by hourly earnings in abstract occupations is shown in blue; hourly earnings in routine occupations divided by hourly earnings in manual occupations is shown in green. Model results are in dashed lines. The series are normalized by their trend value in 1985, where the trend is obtained with an HP filter with parameter 6.25.

of the observed decline in relative hourly earnings.

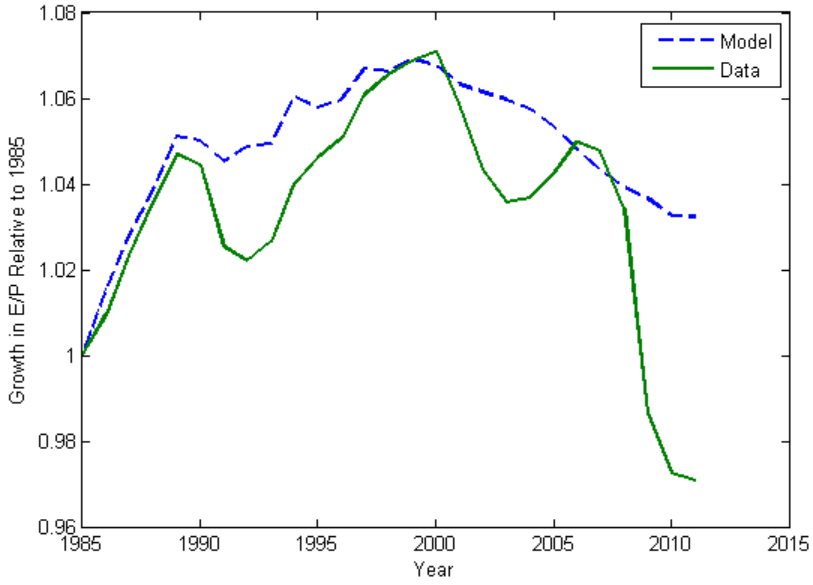
As for hourly earnings in routine occupations relative to manual occupations, the primary difference in the model and the data series occurs in the early 1990s. The model's results show an almost constant ratio of hourly earnings in routine occupations relative to manual occupations, with only a slight negative trend. The data also appears to be relatively trendless (although with possibly a slight positive trend), but only after 1994. Prior to 1994, there is an slight upward trend in this ratio in the data, which is not mirrored in the model. Generally, the model is qualitatively and largely quantitatively consistent with the data along this margin.

Thus, relative to the employment shares results shown before where employment was overstated in abstract occupations and understated in manual occupations, relative hourly earnings to routine occupations is understated for abstract occupations, but quite close to the data for manual occupations. This suggests that reconciling the model to these two features of the data is about more than simply changing the total allocations of workers across occupations.

1.6.2.3 Employment to Population Ratio

The model's predictions regarding the aggregate employment to population ratio are presented in Figure 1.17. Per the calibration of the disutility bounds, the employment to population ratio matches the data well in the period between 1985 and 1992. From 1992 to 2000, the model continues to capture a sizable fraction of the trend movements in the employment to population ratio, which is not surprising as the decline in labor disutility has been indexed to match the movement in female labor force participation. Of greater interest is the model's results regarding the decline in employment from 2000 to 2011, after movement in $\bar{\chi}_t$ has ceased. In this last decade, the model predicts a 3.3% decline in the employment to population ratio, compared to the 9.4% decline observed in the data. That the model is able to account for 35% of the decline of the employment to population ratio is significant, given that the model neither includes the 2000s recession or the Great Recession. This suggests that a modest fraction of the recent decline in the aggregate employment rate may be due to secular factors associated with the ongoing polarization of the labor

Figure 1.17: Employment to Population Ratio, Model vs. Data, 1985 to 2011



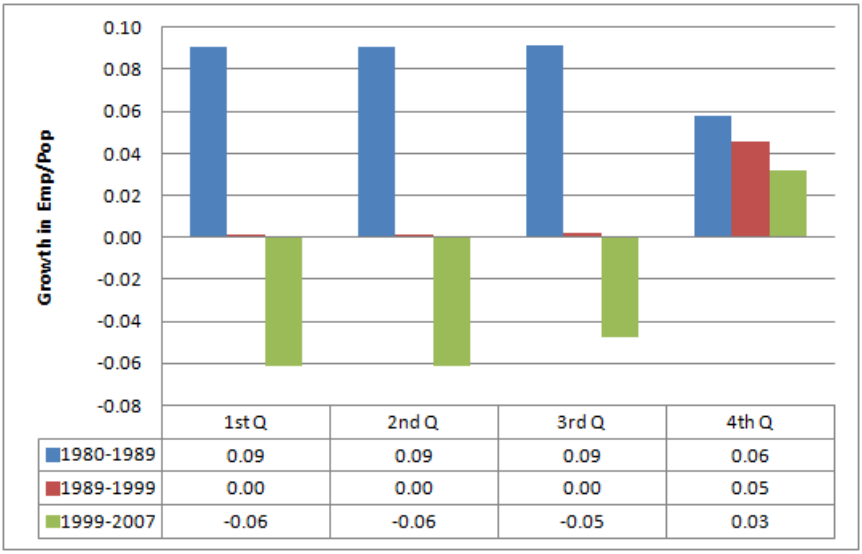
Data on the employment to population ratio is simply the annual average of the series used in Figure 1.6. The employment to population ratio for the model is simply total employment, as the population is constant at 1.

market.

This decline in the aggregate employment to population ratio is generated by disproportionate declines amongst lower skilled individuals. Figure 1.18 presents the model's growth in employment to population ratios across four quartiles of the skill distribution for each of the time periods 1980-1989, 1989-1999, and 1999-2007. As there is no measurable educational margin in my model, I can only consider a qualitative comparison between the model results and what is observed in the data, shown previously in Figures 1.7 and 1.8. In the 1980s, the decline in $\bar{\chi}_t$ increases the employment to population rate substantially for all groups, though not as much in the highest quartile, where many skill levels are already at full employment. However, for the subsequent two decades in the 1990s and 2000s, the model captures the pattern of employment rate declines being greater for lower skilled workers. As higher skilled workers continually transition to abstract occupations and wage growth in abstract occupations in faster than in other occupations, this raises their employment rate relative to lower skilled workers.

The one discrepancy in this qualitative comparison of employment rates across skill levels lies

Figure 1.18: Growth in Employment to Population by Skill Quartile and Decade, Model Outcomes



	1st Q	2nd Q	3rd Q	4th Q
1980-1989	0.09	0.09	0.09	0.06
1989-1999	0.00	0.00	0.00	0.05
1999-2007	-0.06	-0.06	-0.05	0.03

Growth in employment rates per skill quartile is computed using the growth of the integral of the employment rate for each quartile of the skill distribution.

in the nature of changes in the 1980s and 1990s. The effects of a declining $\bar{\chi}_t$ have a near symmetric impact in affecting employment rates across the skill distribution. However, the observed increased employment rates among females in this period, which $\bar{\chi}_t$ is proxying for, seem to show greater signs of skill bias in favor of higher skilled workers.

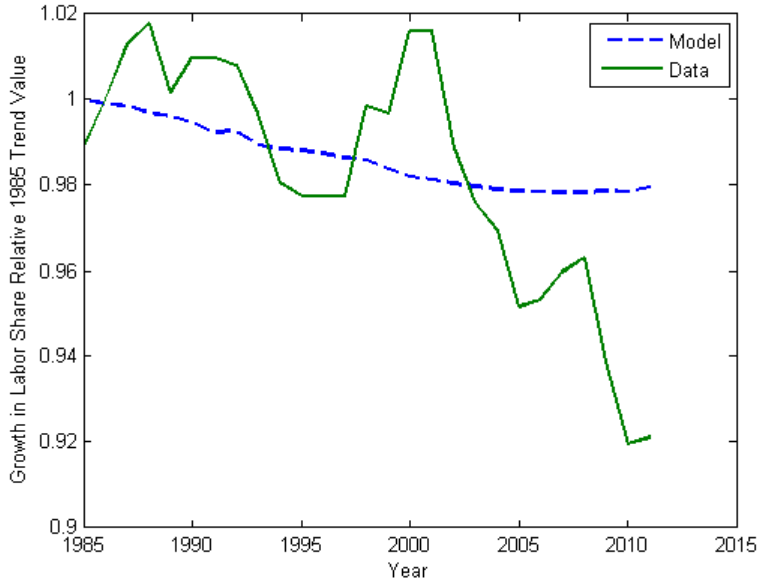
1.6.2.4 Labor's Share of Income

Finally, Figure 1.19 reports the labor share of income for the model and the data. Similar to the hourly earnings series, as there are substantial cyclical fluctuations in the data, I normalize the data and the model results to their trend values in 1985.

The model predicts a 1.9% decline in the labor share between 1985 and 2011, with the decline occurring fairly smoothly throughout the sample until slowing down in the late 2000s. This represents roughly a quarter of the measured decline in the data, where there was an 8.1% decline over the same period.[37] Examining the timing of these declines, the model's results regarding the labor share actually exactly match the data until about the early 2000s. It is only after this point when the

[37]Notably, this is a smaller final amount than found by Karabarbounis and Neiman (2014), who attribute 50% of the labor share decline to investment specific technological change.

Figure 1.19: Labor's Share of Income 1985 to 2011, Model vs. Data

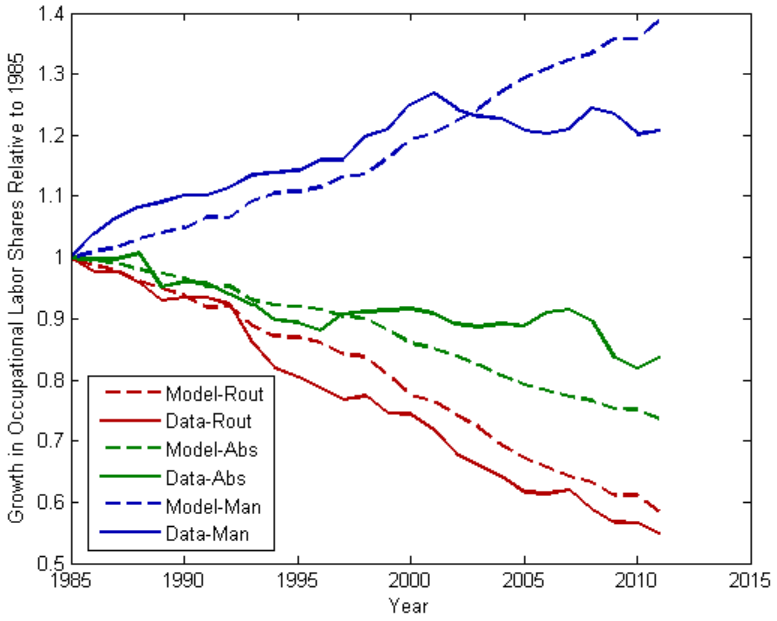


Data on the labor share is simply the annual average of the series used in Figure 1.9; all series are normalized by their trend value at 1985, where the trend is computed with an HP filter with parameter 6.25.

model and data diverge, with the model labor share leveling out while it has continued to decline in the data.

As in Section 2, the labor share can be decomposed into occupational labor shares. The results regarding relative hourly earnings and employment shares can be seen in the evolution of these occupational labor shares, shown in Figure 1.20. The model's predictions regarding occupational labor shares from 1985 to 2000 are fairly close to the data, as with the occupational employment shares. The discrepancies over this early period are larger than with employment shares, however, because the model does not completely replicate the observed decline in relative hourly earnings at this time. The model's results regarding the labor share then diverge from the data just before 2000, reflecting the over and underprediction of abstract and manual employment shares during this period. However, although the model's overpredictions in growth are near symmetric between abstract and manual occupations, the net level effect leads to a higher aggregate labor share, as both employment and hourly earnings in abstract occupations are higher than in manual occupations.

Figure 1.20: Occupational Labor Shares of Total Income 1985 to 2011, Model vs. Data



Occupational labor shares have been constructed as in Figure 1.10. All series are normalized to 1 in the year 1985.

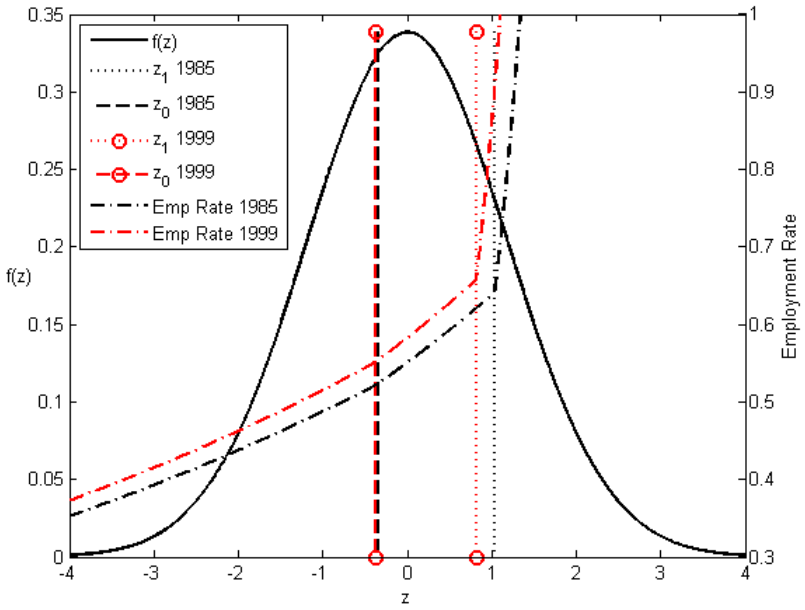
1.6.3 Underlying Mechanisms

Given these results, before reaching any summary conclusions, it is helpful to gain additional insight into how these results are being generated through the model's mechanisms. In particular, I highlight how movements in the employment rate and the occupational sorting thresholds relate to the outcomes observed for employment and hourly earnings.

1.6.3.1 Employment

I show how employment rates and sorting thresholds have changed over time in Figures 1.21 and 1.22, which plot the employment rates and sorting thresholds for 1985 and 1999 and for 1999 and 2011, respectively. In Figure 1.21, we can see that between 1985 and 1999, the only movement in the thresholds for occupational choice occurs for z_1, as higher skilled routine workers shift to employment in abstract occupations. Meanwhile, the employment rate is rising for all skill levels in this time period, through a combination of movements in q_t and $\bar{\chi}_t$. As $\bar{\chi}_t$ decreases, the employ-

Figure 1.21: Employment Rate and Sorting Thresholds, 1985 and 1999



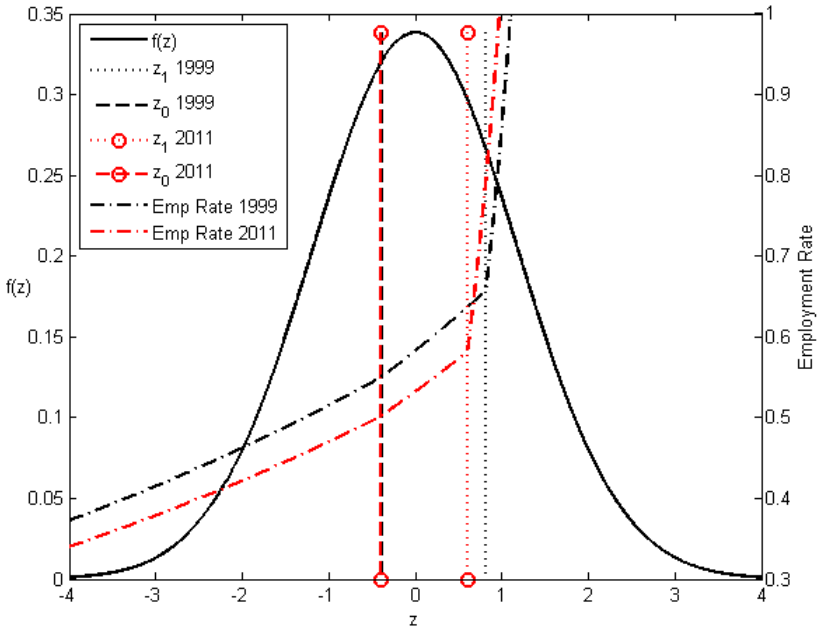
z_1 is in dotted lines; z_0 is in dashed lines. The employment rates are in dash-dotted lines and plotted against the right y-axis. Black represents 1985 and red represents 1999.

ment increases almost uniformly across the skill distribution, with the only nonuniform movements occurring for skill levels already at full employment. However, as q_t increases, this decreases the sorting threshold z_1 and the workers who now shift into abstract occupations observe an higher employment rate, owing to comparatively faster wage growth and the increased productivity in this occupation.

In Figure 1.22, the effect of q_t on the lower end of the skill distribution becomes clear as movements in $\bar{\chi}_t$ cease. Between 1999 and 2011, when the only exogenous change is in q_t, z_1 continues to decline, inducing more workers to enter the abstract occupation and z_0 continues to remain constant. In terms of employment rates, now, without offsetting movements in $\bar{\chi}_t$, the increased wealth of the household reduces the marginal gain from working for all workers, lowering the employment rate for all workers. This is offset somewhat for abstract workers, whose wages are growing faster than in other occupations, but for individuals outside of the abstract occupation, the combination of increased wealth and slow wage growth leads to a substantial decline in employment participation.

This highlights how movements in $\bar{\chi}_t$ have had near symmetric effects across the skill distri-

Figure 1.22: Employment Rate and Sorting Thresholds, 1999 and 2011



z_1 is in dotted lines; z_0 is in dashed lines. The employment rates are in dash-dotted lines and plotted against the right y-axis. Black represents 1999 and red represents 2011.

bution and how ongoing increases in q_t simply increase the share of labor employed in routine occupations.

1.6.3.2 Hourly Earnings

In thinking about the underlying model mechanisms for relative hourly earnings across occupations, it is helpful to think of these results as arising through the intersection of two margins - the relative wages per efficiency unit and the relative efficiency units per worker. For example, consider the relative average hourly earnings between routine and abstract occupations:

$$\frac{w_{rt}N_{rt}}{E_{rt}} / \frac{w_{at}N_{at}}{E_{at}} = \left(\frac{w_{rt}}{w_{at}}\right)\left(\frac{\frac{N_{rt}}{E_{rt}}}{\frac{N_{at}}{E_{at}}}\right)$$

Figure 1.23 plots these two margins for the relative hourly earnings between routine and abstract occupations, with relative wages on the left axis and relative efficiency units per worker on the right axis. Consistent with the occupational cutoff movements shown above, the wage in rou-

Figure 1.23: Relative Wages per Efficiency Unit and Relative Efficiency Units per Worker, Routine vs. Abstract Occupations



Relative wages, w_{rt}/w_{at}, are given by the solid line and plotted against the left y-axis. Relative efficiency units per worker, $\frac{N_{rt}}{E_{rt}}/\frac{N_{at}}{E_{at}}$, are given by the dashed line and are plotted against the right y-axis.

tine occupations relative to abstract occupations has been declining throughout time. However, at the same time, the efficiency units per worker in routine occupations has been increasing relative to abstract occupations. Why has the relative efficiency units per worker been increasing over time in favor of routine occupations? This occurs due to a selection effect - newly transitioning workers into abstract occupations become the least skilled workers in that occupation, providing the fewest efficiency units of all skill types there. The transitioning of these workers dilutes the average skill level in both abstract and routine occupations, but the dilution is greater in abstract occupations, particularly because of the higher productivity of workers in that occupation.

Thus, the important observation for the model here is that increases in employment in abstract occupations are only able to occur through inducing lower skilled workers to switch occupations. This may not be a particularly realistic assumption to make regarding the recent increase in employment in abstract occupations and I briefly address this in my summary.

1.6.4 Summary

Summarizing the above results, I find that

- The model explains employment shares exactly from 1992 to about 2000, but substantially overpredicts the share of labor in abstract occupations and underpredicts the share in manual occupations in the most recent decade. However, the model accounts for the continued decline of routine workers throughout this period.

- Between 35-55% of the decline in routine hourly earnings relative to abstract occupations can be explained in the model. The model's predictions regarding hourly earnings in routine occupations relative to manual ones is qualitatively and quantitatively consistent with the observed data, with the exception of a brief period in the early 1990s.

- The model captures 33% of the observed decline in the employment to population ratio from 2000 to 2011 and is consistent in many ways with the qualitative results regarding employment to population ratios across the skill distribution. However, proxying for female labor force participation changes using $\bar{\chi}_t$ appears to overstate the changes in participation for lower skill levels.

- 25% of the labor share decline is accounted for by the model. However, the model accounts for almost all the decline in the labor share until the year 2000, but thereafter is unable to match the observed decline. This appears to be partially driven by the model's overprediction of abstract employment in this decade.

Given these results, I conclude that technological change can explain a significant fraction of the observed changes in occupational employment shares and relative hourly earnings, but this explanation seems less well suited to capturing labor market outcomes since the year 2000. In particular, the non-linear behavior of employment shares in abstract and manual occupations and these consequences for the labor share of income remain puzzles given the continuing rapid progress of

technological change throughout this time period. However, it still is able to explain a modest fraction of movements in the aggregate employment to population ratio in this period.

A valuable question to consider is how much of what the model is unable to explain quantitatively is potentially related to assumptions made in the modeling strategy. The analysis of underlying model mechanisms shown above highlights two particular aspects of the model which might be valuable to expand upon in future refinements to this framework: proxying for female labor force changes with $\bar{\chi}_t$ and the assumption of a static skill distribution. The modeled decline of the disutility of work implies a near symmetric effect across skill levels of inducing workers to enter the labor force. However, given that female labor force participation changes in the 1980s and 1990s have not been entirely symmetric, this suggests a more detailed approach to this may be prudent. This could be particularly important, given that the calibration strategy targets employment shares, whose movements in the calibration period could be partially driven by movements in differential labor force participation changes across the skill distribution. For example, if some of the observed change in the abstract employment levels is due to differential participation changes across skill levels, the model might overestimate the degree of complementarity between abstract work and routine tasks.

Second, throughout this analysis, I have maintained a benchmark of a constant skill distribution. However, there are well known changes in educational attainment over the past several decades which might suggest changes in the skill distribution as modeled here.[38] Again, if supply factors are partially driving the changes in abstract employment shares in the 1980s when the model calibrates to match employment shares, the complementarity of abstract tasks might be overstated. Further, changes in the supply of skill could potentially have strong consequences for the observed skill dilution effect prominent in the relative hourly earnings between abstract and routine workers. If there are increases in the skill level, then some of the increased demand for abstract tasks may be met by a greater supply of high skilled workers, and not from lower skilled workers transitioning into this occupation. This could also explain some of the brief rise in the 1990s in

[38] Acemoglu and Autor (2011) provide a summary of some of these changes.

the hourly earnings of routine workers relative to abstract workers. An important extension to this quantitative framework would be to incorporate some measure of a changing skill distribution over time, be it through exogenous movements in the distribution of skill or through endogenous skill accumulation.

1.7 Conclusion

Much of the discussion surrounding the polarization of the labor market and recent changes in aggregate labor market outcomes has centered around the role of technological change. In this chapter, I take a first step towards a formal quantitative evaluation of the role of technological change in explaining these phenomena by augmenting the standard growth model with an occupational tasks framework in production, a Roy model for occupational choice, and equipment investment specific technical change. I further attempt to reconcile the lack of polarization across employment rates by including an explicit employment versus non-employment decision margin.

With this model, I show that the model is able to explain a substantial fraction of these changes through the year 2000, but is less effective in explaining a number of changes occurring since 2000. The model also explains a modest fraction of the recent decline of the employment to population ratio, suggesting a potentially important role for secular forces in explaining these recent changes. And this decline is consistent qualitatively with evidence regarding employment rate differences by skill. The model is additionally able to explain 25% of the decline in the labor share from 1985 to 2011.

While it is too early to reach a final conclusion regarding the total effects of technological change, this initial evaluation suggests a substantial role for it in the recent polarization of the labor market and other changes in aggregate labor market outcomes, but primarily through the year 2000. These results have also highlighted several margins to refine in future work, particularly the labor supply margins capturing changes in female labor force participation and the evolution of the skill distribution over time, and further refinements are necessary to evaluate how sensitive these

results are to the assumptions made here.

But a nagging puzzle which may not be immediately resolved by such refinements is this: why has growth in the abstract occupational employment share slowed so much recently? After growing steadily for several decades, it suddenly slowed in the 2000s, with faster growth occurring in manual, lower-skilled occupations. Autor and Dorn (2013) have given some attention to the growth of low skill service occupations in this time period, and Beaudry et al. (2013) have begun to consider possible declines in labor demand for high skilled labor since 2000, but there has been little work done on understanding this most recent change in occupational employment dynamics. I consider this an important area for future research.

1.8 Appendix

1.8.1 Data Sources and Construction

For data on changes in the distribution of occupational employment and hourly earnings, I use both the CPS Basic Monthly files as well as the Annual March CPS supplement files. All data evidence on employment figures is drawn from the monthly files; all data evidence regarding hourly earnings is from the March supplement. Additionally, for the construction of the relative price of equipment investment to consumption, I use detailed price series on consumption and investment from the National Income and Product Accounts. Before describing the procedures for preparing the data, however, I consider the industry and occupational code systems obtained from Autor and Dorn (2013).

1.8.1.1 Occupational and Industry Classifications

For time-consistent occupational and industry classifications, I use the modified occupational classification system used in Autor and Dorn (2013) and the modified industry classification system used in Autor, Dorn and Hanson (2013a). However, the occupational classification system does not account for the recent 2010 revision to the census occupational codes, and the industry classification system does not extend beyond the 1980 industry codes. Thus, I use crosswalks from the census bureau to extend these classifications systems forward and backward as needed. A full concordance of this extension method is available upon request.

I make one other modification to the above mentioned classification systems. While these classification systems generally do an excellent job of providing time consistent measures of occupations, there are still subtle discontinuities each time the classification systems are revised. To eliminate these discontinuities, I restrict the share of employment in a given occupational class remain constant across the redefinition. As this has generally mild effects with the CPS monthly files at the level of aggregation I consider, I prefer the use of these data for more precise measures of employment shares. But even in the annual data, imposing this restriction does not greatly alter

the conclusions of the data.

To map occupations into the three task classes - abstract, routine and manual - I use the mapping from Autor and Dorn (2013) which looks at task intensity across occupations and determines which occupations correspond to which classes based on whether or not the task intensity scores are above average. With this mapping, abstract, routine and manual occupations are identified as follows:

Abstract: management, professional, technical, financial sales and public safety occupations

Routine: precision production and craft occupations, machine operators, assemblers and inspectors occupations, and clerical/administrative support and retail sales occupations.

Manual: transportation, construction, mechanics, mining occupations and low skill service occupations

It is worthwhile to note that there have been a wide variety of mappings employed to relate occupations to tasks. Alternate classification systems have often additionally distinguished jobs by two categories: routine vs. non-routine and cognitive vs. manual.[39] As in Autor and Dorn (2013), abstract occupations are essentially non-routine cognitive jobs and manual jobs are non-routine manual jobs, where as routine jobs are either routine cognitive or routine manual. An often difficult to categorize group of jobs are mechanic/repair/construction/mining occupations, as they have both strong manual and routine content. I follow Autor and Dorn (2013) and label these as manual jobs, although Jaimovich and Siu (2012) consider these to be routine occupations as well, and Smith (2013) labels these occupations as "other."

1.8.1.2 CPS Basic Monthly Files

For data on employment trends, I use the CPS basic monthly files from January 1976 through to July 2013 on employed individuals between the ages of 16 and 65 (though all results are robust to removing the age ceiling). I restrict attention to employment in the non-farm business sector, and thus drop workers if they are employed in the agriculture or public administration sectors, or if they are not employed in the private sector. For all time series constructed, I filter them with a

[39] Acemoglu and Autor (2011) use this particular classification scheme.

band pass filter, removing variation at frequencies between 2 and 18 months, to remove seasonal "noise."

1.8.1.3 CPS March Supplement Files

For hourly earnings data, I use the Annual March CPS Supplement files from 1976 to 2012, obtained through IPUMS (King et al. (2010)), which reports information on individuals' income for the prior year. I similarly restrict attention to workers employed in the non-farm business sector, and thus drop workers if they are employed outside of this sector. My focus on average hourly earnings is generally consistent with the recent literature.[40] It also allows me to construct aggregate factor shares for occupational labor income.

To construct average hourly earnings, I need accurate measures of hours and income, both of which have potential issues in the March CPS files. For hours non-respondents, I impute hours worked as in Krusell et al. (2000). For income, there are a variety of potential measurement issues, the most prominent being the top-coding that is applied in the data. Income recorded in the CPS is subject to two levels of top-coding - public use data top-coding and internal top-coding. In the data prior to the survey year 1996, income levels exceeding the public top code are simply listed at the top code value. From 1996 onwards, income levels at the top code are replaced with an average of actual income values for individuals with similar characteristics. Larrimore et al. (2008) provide comparable measures for this "cell-mean" procedure dating back to the 1976 survey, so I use these cell-means until the CPS begins using them in practice. However, this procedure does not provide any adjustment to compensate for the internal top-coding. To make this correction, I fit a Pareto distribution to the tail of the income distribution, using the 90th and 95th percentiles, and use it to impute a correction term to the cell-means to compensate for lumpy changes induced by the internal top code.

To obtain a comparable aggregate measure of the labor share from the household survey, I need

[40] See Footnote 8 in Acemoglu and Autor (2011) for a more in-depth discussion of using hourly earnings or average wages.

to impute earnings of the self-employed. I follow the BLS procedure of imputing self-employed wage earnings as the average wage (over all observations where the implied wage is greater than half the minimum wage) multiplied by the usual number of hours worked. In cases where individuals both both business income and wage income, I imputed self-employment earnings as the average wage times either 2/3 the total hours worked if the primary job was self-employment, and 1/3 the total hours worked if the primary job was wage and salary employment. All results are robust to varied modifications in this imputation procedure.

Even with the imputed self-employed income, the aggregate income series derived from the March supplement files will be unlikely to completely match the aggregate income series the BLS calculates from the establishment survey. This is because income in the establishment survey also includes employer contributions to health care, as well as payments in kind (particularly, the exercise of non-qualified stock options). As a result, the estimated aggregate income series from the household is always short about 80-85% of the total income reported by the BLS; however, despite the level gap, the dynamics of the two series match very closely. As such, to match the aggregate labor share, I evenly divide the missing fraction amongst the income series for each occupational group.

The final two adjustments needed for the March supplement data are with regard to the survey redesigns that have occurred in the past several years. In particular, changes in how the CPS was administered in 1994 and 1995 dramatically affected measures of income as the rate of non-response changed substantially and asymmetrically. As a result, there are some unusual discontinuities in the aggregate income series. I thus impose the restriction that the growth rates of aggregate income generated from the CPS for the years 1993-1995 match the growth rates of the BLS measure of aggregate income to smooth out the aggregate income measure obtained from the CPS data. The other adjustment that is made, though is ultimately not relevant for the final results, is to make a correction for the changing in the weighting scheme implemented in 2003. Aggregate statistics reported by the BLS adjust the population weights to shift this discontinuity back in time to 2000; I employ a comparable procedure to redistribute population weights consistently between

2000-2003.

With the now adjusted measure of aggregate labor income, I compute average hourly earnings as the ratio of aggregate labor income to aggregate hours worked.

1.8.1.4 Relative Price Data

As discussed in the text, I use the quality adjusted price series for the price of capital equipment investment, the same as used in DiCecio (2009). For data on the price of consumption, I use price indices for Personal Consumption Expenditures and construct a Tornqvist index for the price of non-durable consumption and non-housing, non-energy services, as in Cummins and Violante (2002). I also go one step further and correct for price measurement bias in the consumption prices. Using the recommendations of Boskin et al. (1996) and Gordon (2006), I apply corrections for quality bias as estimated in these papers to obtain a more accurate measure of the price of consumption. These corrections require delicacy to apply, as the original findings on consumption price quality bias are for the CPI and not the PCE. But a large fraction of the underlying data in the PCE comes from the CPI and by mapping different consumption subcategories between the underlying data for each of these, I can adjust the growth rate of prices in the consumption subcategories where quality bias is identified. Further details of this mapping/matching procedure and the exact quantitative adjustments made are available upon request from the author. However, all results in the chapter are robust to using the BEA series without bias adjustment.

1.8.2 Industry Robustness Checks for the Measured Decline of Routine Work

It might be the case that the measured decline of routine employment and hourly earnings is simply capturing changes in the industrial composition of the economy or changes within a specific set of industries, particularly the manufacturing industry. Here I demonstrate that this is not the case.

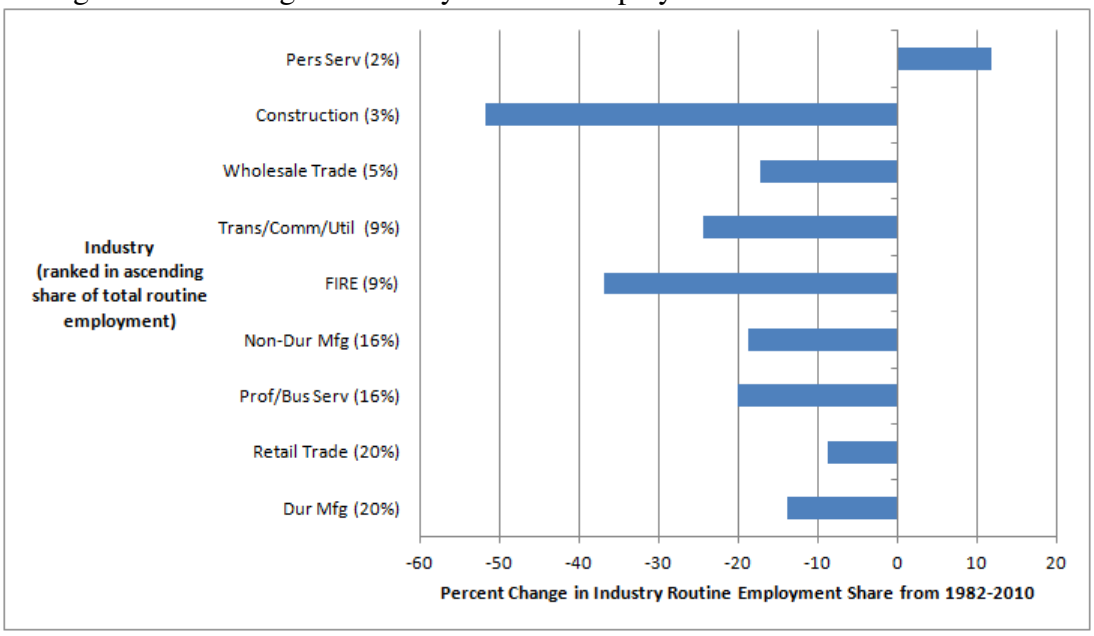
The aggregate routine employment share can be decomposed at the industry level as follows:

$$\frac{E_t^{rout}}{E_t} = \sum_{j=1}^{N} \frac{E_{jt}^{rout}}{E_{jt}} \frac{E_{jt}}{E_t}$$

where j stands for industry. Thus, the routine employment share is written as the share of routine employment within each industry and each industry's share of total employment. Holding the industry shares of total employment fixed at the year 1982, I consider how much of the change in the routine employment share can be explained by only variation in within industry routine employment shares. Performing this exercise, only allowing for within industry changes, I am still able to explain 68% of the decline in routine employment. This suggests that while there may be a role for between industry movements in explaining the decline of routine work, they contribute less than a third of the total change.

Figure 1.24 presents evidence regarding the distribution of routine employment across industries as well as the change in the share of routine work within each of these industries. Industries are presented in increasing order of the share of total routine employment hired in this industry (number of routine workers in industry i divided by the total number of routine workers). A substantial fraction of routine employment is concentrated in manufacturing, but by no means a majority. 36% of all routine employment is in the manufacturing industry, fairly evenly split between durable and non-durable manufacturing. Routine employment is also particularly present in both the retail trade and professional and business services industries, an additional 36% of routine employment. The remaining roughly third of routine employment is divided across all other industries, and is particularly scarce in personal/social services and construction. Perhaps more importantly, however, is that the share of routine employment in total industry employment (number of routine workers in industry i divided by the total number of workers in industry i) has been substantially declining in almost every industry. The largest reductions in industry shares of routine work have been in construction and FIRE, but most industries have observed at least a 10% reduction in their share of routine employment. Thus, employment in routine occupations is spread out across a number of industries, not just manufacturing, and routine employment has been declining in virtually every industry.

Figure 1.24: Change in Industry Routine Employment Shares from 1982-2010



Data comes from the CPS basic monthly files. Occupational and industry definitions come from Autor and Dorn (2013). Annual employment shares are obtained by averaging monthly employment shares. The percentage change in industry routine employment share for industry i is given by $\Delta ln(\frac{Emp_{it}^{routine}}{Emp_{it}^{total}})$. Results are robust to varying the start and end dates of the sample. Additional details regarding data construction and sources are available in the Appendix.

Table 1.2: Changes in Relative Occupational Hourly Earnings Across Industries

	Routine/Abstract			Routine/Manual	
	1982-2011	1992-2011		1982-2011	1992-2011
Pers. Serv.	-28%	-11%		-18%	-3%
Construction	8%	-5%		13%	-9%
Wholesale Trade	-19%	-7%		14%	16%
Trans/Comm/Util	-5%	-8%		2%	-7%
FIRE	-33%	-25%		3%	5%
Non-dur Mfg	-20%	-24%		-1%	-6%
Prof/Bus Serv	-7%	-2%		9%	5%
Retail Trade	-11%	-4%		20%	11%
Dur Mfg	-21%	-16%		4%	-4%
TOTAL	-17%	-12%		4%	-2%

Data is from the CPS March Supplement files on employed persons in the non-farm business sector. Industry classifications come from Autor, Dorn and Hanson (2013a). Average hourly earnings are constructed by dividing total income for an occupational group by total hours worked by individuals in that occupation. Relative hourly earnings is normalized to 1 in 1975. Industries are listed in increasing share of total routine employment. Additional details regarding data construction and sources are available in the Appendix.

Table 1.2 presents changes in the relative average hourly earnings between occupations across industries for both the period of 1982-2011 and 1992-2011. Similar to the employment growth shares, we see that the relative hourly earnings between routine and abstract workers has declined fairly uniformly declined across all occupations, although in some occupations more than others. As for the relative earnings between routine and manual workers, although there was some growth in the 1980s, we see that most industries see a decline and/or negligible change in this relative hourly earnings since the 1990s. The exception here is the trade industry, as there have been significant gains in the relative hourly earnings of routine workers to manual workers in both wholesale and retail trade. This exception aside, the patterns in individual industries are all broadly consistent with the aggregate, suggesting that the decline in routine work is not limited to a small subset of industries.

1.8.3 Solution and Calibration Methods

1.8.3.1 Solution Method

As mentioned in the text, I use a multiple shooting algorithm to solve for the transition path between two steady states of the model economy. The exact solution algorithm is as follows:

1. Given values for the model parameters, solve for the initial and terminal steady states.

2. Use the initial steady state value of consumption as a guess for the value of consumption in the first period and set upper and lower bounds around this guess (typically 2.5% in either direction).

3. Given the initial guess value for consumption, capital stocks from the initial steady state, and the time paths of the shocks, solve the model for period 1 and then iteratively solve the model forward along the entire time path of shocks and then an extended period of N periods where the shocks are held constant at their terminal values.

4. Check the value of consumption at the end of the path and the extended N periods and compare it to the terminal steady state value of consumption. If consumption is too high, update the guess by choosing the midpoint of the lower bound and the current guess; similarly if consumption is too low (standard monotonicity argument). Employ this bisection algorithm until the difference in guesses in adjacent periods has converged.

5. Now given a solution for the first period's consumption, guess a value for the second period's consumption given the capital investment decisions from period 1, and solve the time path again, this time allowing for $N+1$ periods of the shocks being held at their terminal values (Always hold the number of periods the model is solved forward constant, so with each successive iteration, add an additional extended period with shocks at their terminal values.). Obtain a value for consumption and capital investment decisions in period 2.

6. Perform this solution routine to obtain a time series for consumption and capital stocks for the requisite number of periods.

In theory, a simple single shooting algorithm could work, and the iterative procedure would not be needed, but as is well known, single shooting is known to be unstable in explosive saddle-point systems and accumulating numerical errors may creep in with successive iterations that make convergence near impossible (see Lipton et al. (1983) for a discussion of this and multiple shooting in economics). In addition to stability, the further advantage of an iterative procedure such as this is manifest in calibration when the targets in calibration are changes along the transition path. To calibrate the model, only the number of periods needed for the calibration need to be solved, and not the entire time path, thus reducing computational time.

1.8.3.2 Calibration Approach

Steady State

Using the model's steady state, I am able to calibrate most parameters; the values for the Cobb-Douglas share of structures (α) and the depreciation rates (δ_s,δ_e) can be obtained without even numerically solving the initial steady state.

As for the parameter α, along a perfect foresight path, ex post returns to capital structures and capital equipment are equated as follows:

$$(1-\delta_s)+r_{st+1}=\frac{q_t}{q_{t+1}}(1-\delta_e)+q_t r_{et+1}$$

Additionally, the constant returns to scale nature of the production function implies:

$$\frac{r_{st}K_{st}+r_{et}K_{et}}{Y_t}=1-Lsh_t$$

where Lsh_t is the total labor share of income at date t. Given data on the labor share, nominal output, the time path of q_t, and values for the depreciation rates, these equations yield two equations in two unknowns and can be solved to find time paths of the rental rates in the data. I then construct α as:

$$\alpha = \overline{\left(\frac{r_{st} K_{st}}{Y_t} \right)}$$

where the average is taken over the period 1980 to 2011.

Data for capital equipment and capital structures comes from the BEA data on Fixed Assets for the non-farm private sector. For structures capital, real quantities have been constructed using BEA price deflators; for equipment, following the procedure set forth in Cummins and Violante (2002) using the perpetual inventory method with the quality adjusted prices for capital equipment.

Depreciation rates are calibrated similarly, as in steady state, $\delta_s = \frac{I_s}{K_s}$ and $\delta_e = \frac{qI_e}{K_e}$; thus, for both depreciation rates, I take the average of these (effective) investment rates over the time period 1980-2011 and use these as the values for δ_s and δ_e.

The parameters to calibrate by numerically solving the initial steady state are: the share parameters in the CES nests (μ_j), the initial support for the distribution of the disutility of labor ($\bar{\chi}_{1980}$), the productivity parameters in occupational specific productivities (a_j), and the variance of the skill distribution (σ_z^2).

The share parameters are calibrated numerically to match the initial factor shares for different occupational types of labor.[41] In practice, guessing a variety of different values for these share parameters can create difficulty in finding the initial steady state if the initial guess is not a good one. Thus, instead of directly guessing the three share parameters, I make a guess regarding the occupational skill sorting thresholds (z_0, z_1), and then given these values, can solve for the share parameters that produce the proper initial factor shares simply using the production side of the model. Successful calibration occurs when the guesses regarding occupational skill sorting thresholds converge to the actual values in the initial steady state.

The remaining parameters to calibrate then are those regarding the skill distribution, disutility

[41] As income is only available at an annual frequency in the CPS, but the most accurate measure of occupational employment shares comes from the monthly CPS files, I slightly modify the annual files to be consistent with the implied employment distribution by multiplying factor shares by the relative employment shares across the monthly and annual files. Results are robust to this procedure.

distribution and occupational productivities. The initial disutility distribution parameter ($\bar{\chi}_{1980}$) is calibrated to match the initial employment to population rate. Two of the occupational productivity parameters - those for routine and abstract labor - are calibrated to match the initial employment distribution in 1980, and the final parameter for manual productivity is calibrated to match the initial relative variance of log incomes between routine and manual employed workers. The variance of the skill distribution is similarly calibrated to the variance of log income in abstract occupations (though could easily be calibrated to match the variance of log income for all workers). These last calibrations require computing the model's variance over log income for different occupational types. Given the Normal distribution for skill, these variances can be solved analytically. For example, the cross sectional variance of log income for abstract workers at time t is given by:

$$
\begin{aligned}
Var\left[ln(W_t(z)) \mid z > z_{1t}, \chi < \hat{\chi}(z)\right] &= Var\left[ln(w_a) + a_a z \mid z > z_{1t}, \chi < \hat{\chi}(z)\right] \\
&= a_a^2 Var\left[z \mid z > z_{1t}, \chi < \hat{\chi}(z)\right] \\
&= a_a^2 \mathbb{E}\left[z^2 \mid z > z_{1t}, \chi < \hat{\chi}(z)\right] - \\
&\quad a_a^2 \left(\mathbb{E}\left[z \mid z > z_{1t}, \chi < \hat{\chi}(z)\right]\right)^2 \\
&= a_a^2 \frac{1}{E_{at}} \int_{z_{1t}}^{\infty} \int_0^{\hat{\chi}(z)} z^2 \frac{d\chi}{\bar{\chi}} dF(z) - \\
&\quad a_a^2 \left(\frac{1}{E_{at}} \int_{z_{1t}}^{\infty} \int_0^{\hat{\chi}(z)} z \frac{d\chi}{\bar{\chi}} dF(z)\right)^2 \\
&= a_a^2 \frac{1}{E_{at}} \int_{z_{1t}}^{\infty} z^2 \frac{\hat{\chi}(z)}{\bar{\chi}} dF(z) - \\
&\quad a_a^2 \left(\frac{1}{E_{at}} \int_{z_{1t}}^{\infty} z \frac{\hat{\chi}(z)}{\bar{\chi}} dF(z)\right)^2
\end{aligned}
$$

$$= a_a^2 \frac{1}{E_{at}} \left(\int_{\hat{z}_{at}}^{\infty} z^2 dF(z) + \left(\frac{w_{at}}{C_t \bar{\chi}^{\nu}} \right)^{\frac{1}{\nu}} \int_{z_{1t}}^{\hat{z}_{at}} z^2 e^{a_a z} dF(z) \right) -$$
$$a_a^2 \left[\frac{1}{E_{at}} \left(\int_{\hat{z}_{at}}^{\infty} z dF(z) + \left(\frac{w_{at}}{C_t \bar{\chi}^{\nu}} \right)^{\frac{1}{\nu}} \int_{z_{1t}}^{\hat{z}_{at}} z e^{a_a z} dF(z) \right) \right]^2$$

where I have used E_{at} to denote the number of employed workers in abstract workers at time t and $\hat{z}_{at}$ to denote the threshold skill level above which individuals always choose to work. Given this form, and values for the model's steady state, it is straightforward to complete the computation using a completing the square argument and standard facts regarding moments of truncated normals. Similar expressions go through for the other occupational groups as well.

Along the Transition Path

There are four parameters to be calibrated along the transition path - the three elasticity of substitution parameters (γ,ρ,σ) and the terminal value for the time path of the support for the disutility distribution ($\bar{\chi}_{term}$). Before I discuss the exact calibration methods, I first give a little more detail on the construction for a time path for the upper bound of the disutility distribution.

I determine the time path of the disutility distribution to be as follows. I take data on the annual female labor force participation rate from the FRED database at the Federal Reserve Bank in St. Louis from 1980 to 1999, the year when female labor force participation reached its maximum in the postwar period. Knowing that a decrease in $\bar{\chi}$ will increase the labor force participation, I take the growth rate of the reciprocal of the female labor force participation rate as my measure of the growth rate of $\bar{\chi}$ along the time path and then scale the period by period growth rate by a constant factor so as to reach $\bar{\chi}_{term}$ in the year 1999. From 1999 onward, the value of $\bar{\chi}$ is fixed at $\bar{\chi}_{term}$.

To calibrate along the transition path, I simply solve the values for the employment shares, the employment to population ratio and capital equipment to worker from the steady state until the year 1995. As there is negligible business cycle variation in the relative price series and the female

labor force participation rate, I smooth the targets and the model moments using a band pass filter to remove business cycle frequencies from 1985-1995 and thus calibrate to match the change in the smoothed series between 1985 and 1992.

Chapter 2

Occupational Tasks and the Decline of the U.S. Labor Share

2.1 Introduction

In the past several decades, labor's share of income has been declining. This is a puzzle, as since Kaldor (1961), modern macroeconomics has generally operated under the assumption of a constant aggregate labor share. Two theories have recently arisen trying to explain this recent change. Karabarbounis and Neiman (2014) argue that the rapidly falling relative price of investment, representing technological progress in information and computing technologies, has led to an aggregate shift away from labor toward capital, and thus a decline in labor shares worldwide. In contrast, Elsby et al. (2013) argue against technological change as an explanation of the decline of the US labor share, suggesting instead that an increase in offshoring accounts for the recent decline.

In this chapter, I contribute to the ongoing discussion of the decline of the labor share by adopting a task-based perspective on changes in labor compensation as a lens for viewing the recent decline of the labor share in the United States. In particular, I build on the work of Autor and Dorn (2013) and consider labor employed in occupations specializing in performing one of three different types of tasks - abstract, routine and manual tasks. The recent literature on tasks

and inequality, summarized in Acemoglu and Autor (2011), has identified that workers employed in occupations performing routine tasks have seen comparatively slower wage and employment growth in the past several decades. The common explanations for this recent labor market polarization have been that labor employed in routine tasks can be easily replaced by technology or is highly exposed to import competition or offshoring. As the two leading theories for polarization are the same as the leading theories for the decline of the labor share, an interesting question is whether the recent decline in the labor share is related to the polarization of the labor market.

The first contribution of this chapter is to argue that the answer to this question is yes. I present a variety of evidence to suggest that there is a connection between the decline of routine work and the decline of the labor share. I begin by dividing compensation by occupations and show that the share of output paid as compensation to workers in routine occupations has been falling over time and that this decline has been greater than in all other occupations both in the long run and since the year 2000. Further, I show that this evidence is robust to controlling for the changing industry composition of the economy and that the greatest declines in the routine compensation share of output have been in those industries which have experienced the greatest labor share declines. Finally, I document a robust correlation between measures of routine task intensity and changes in the share of compensation paid to routine workers and declines in labor shares at the industry. All this evidence combines to suggest that the aggregate labor share decline is closely related to the polarization of the labor market and that theories for the decline of the labor share should give close attention to how the income formerly paid to routine workers is being displaced by outside factors such as technology or trade.

The second contribution of this chapter is to present an additional force that has contributed to the decline of the labor share, potentially reconciling the timing challenges faced by these existing theories of the labor share decline. As will be shown later, most of the recent labor share decline is concentrated in the most recent decade, following the year 2000. Elsby et al. (2013) highlight how this presents a problem for a theory such as technological change measured through the relative price of investment, as the relative price has been growing significantly since the 1980s, and has

actually slowed some in the recent decade. Further, standard implications of such a theory through the lens of a neoclassical model would suggest that there should have been corresponding increases in output per hour and capital deepening if technological change has been driving the labor share decline since the 2000s, and neither has been evidenced in this time period. Many of these criticisms hold as well for theories of trade or offshoring which posit that routine employment has been substituted for by labor abroad, as these too would suggest capital deepening and many changes in trade and offshoring have been occurring for some time prior to the turn of the century.

Viewing the labor share decline through the lens of changing compensation patterns in occupational tasks, I observe that in the year 2000 there is a sudden slowdown in the share of output paid to workers employed in high-skilled abstract task occupations. In particular, I argue that this is an important cause for the increased labor share decline in the past decade. Although the share of output in routine task labor compensation has been falling over time, this has been largely offset by increased compensation in abstract task labor compensation through the year 2000, leading to only a minimal decline in the labor share during this period. However, after the turn of the century, the stagnation of growth in the abstract compensation share has removed this offsetting force, leading to a substantial recent decline in the labor share. I document this aspect of the data and observe that it occurs asymmetrically across industries, and is particularly concentrated in trade, information, professional and business service, and education/health industries. Further, the decline is particularly concentrated in managerial, computer science, medical and sales occupations.

While identifying an exact cause for the slowdown in the abstract compensation share is beyond the scope of this chapter, I consider a few possible explanations for this slowdown. In particular, I consider changes in the composition of compensation which may had lead to compensation occurring through poorly measured channels, such as stock options and stock grants. While there is limited evidence available to consider such a hypothesis, the existing evidence on managerial compensation at S&P 1500 firms suggests that this is not likely to explain well the recent slowdown. An alternate explanation is that there has been a slowdown in the demand for high-skilled labor, associated with the IT revolution reaching "maturity" in the year 2000, a theory put forth in

Beaudry et al. (2013). I briefly discuss this theory and its limitations and potential for explaining some of these changes and future work that can be done to improve our understanding of the recent slowdown.

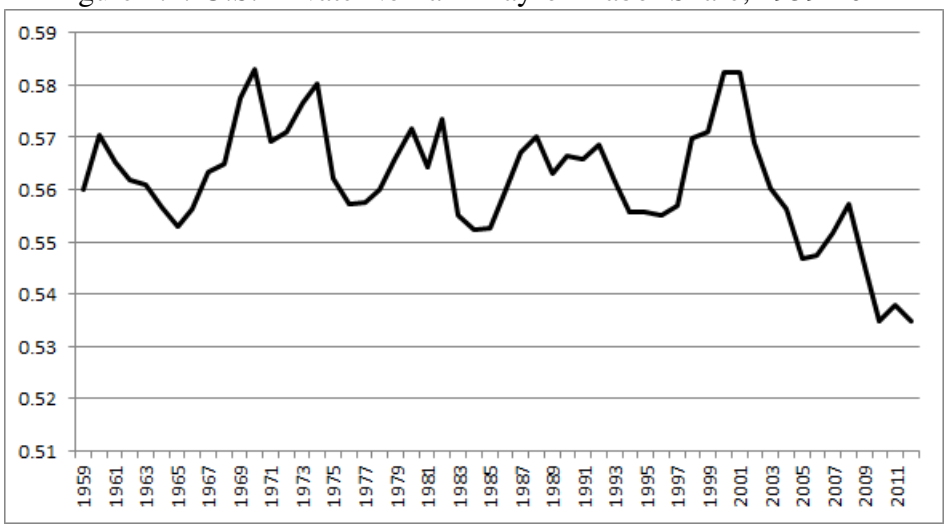
The rest of the chapter is outlined as follows. Section 2 reviews evidence for the labor share decline, making measurement choices so as to be comparable to the data used to divide labor compensation across occupations. Section 3 presents the basic evidence for the relationship between polarization and the labor share decline, as well as the slowdown of compensation in high skilled abstract occupations and the recent acceleration of the labor share decline. Section 4 presents evidence to help further understand the recent slowdown of income paid to workers in abstract occupations, and Section 5 concludes.

2.2 The Decline of the Labor Share

In this section, I document the decline of the labor share in the United States over the past several decades, giving especial attention to the acceleration of the decline since the turn of the century. As the focus of this chapter is to use a occupational tasks framework to understand this decline, it is important to start with a measure of the labor share which can potentially be decomposed into labor shares across different occupations, using data from the Current Population Survey. As such, I focus on the payroll labor share for the private nonfarm sector. In so doing, I closely relate to the measure utilized by Elsby et al. (2013) for much of their analysis. The only difference between their measure and mine is that I do not remove non-profit institutions serving households, but I do remove government enterprises - this is because detailed information on these categories is not available from the CPS and more detail on these nuances is available in the Appendix. Of note, by focusing on the payroll labor share, I do not take any stance on the appropriate imputation for labor income for proprietors, on which there has been no clear consensus.[1] This emphasis on the payroll

[1]Elsby et al. (2013) observe that differences in this imputation can have a substantial effect, and suggest that up to a third of the labor share decline measured by the Labor Productivity and Costs series of the BLS is attributable to their choice of imputation.

Figure 2.1: U.S. Private Nonfarm Payroll Labor Share, 1959-2012



Data is obtained from the National Income and Product Account Tables provided by the BEA. Construction is described in the body of the text.

share for this exercise is further motivated by the documented underreporting of self employed income in household surveys such as the CPS (see Hurst, Li and Pugsley (2014)).

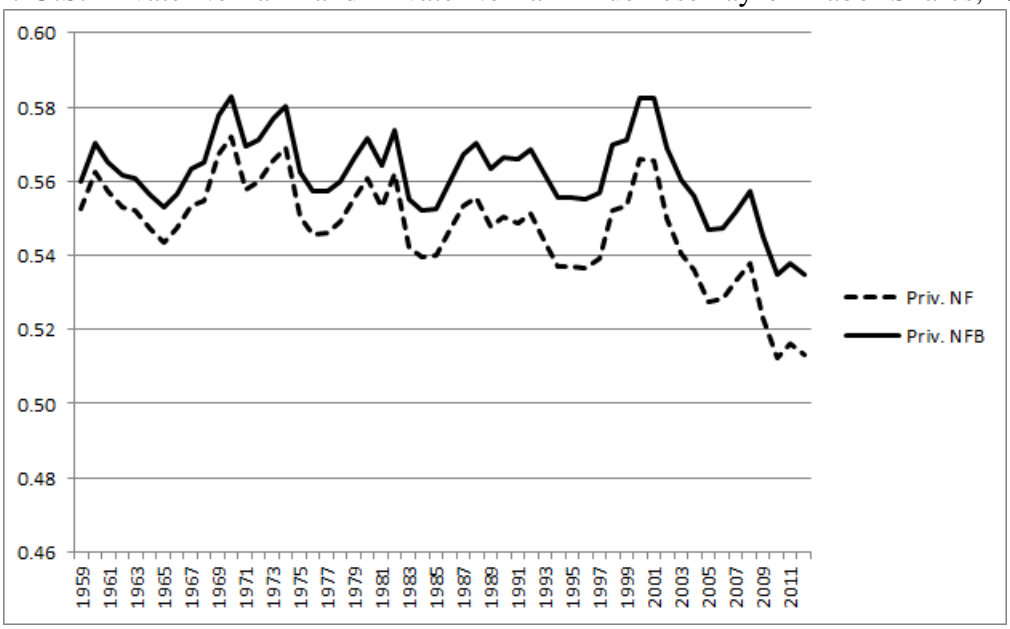
I measure compensation for the private nonfarm sector by subtracting compensation in farm industries and private households from compensation in private industries, using Tables 1.13 and 6.2 of the NIPA.[2] Value added is measured by the sum of value added in the nonfarm business and nonprofits serving households sectors minus value added in government enterprises, using Tables 1.3.5 and 1.13 of the NIPA. The payroll labor share is then given by the ratio of compensation to value added and is plotted in Figure 2.1 for the period 1959 to 2012.[3] For much of this period, the payroll share has been relatively constant, hovering around 57%. However, following a peak in the early 2000s, the labor share fell steadily to a sample low in the year 2005, and has since fallen to nearly 53% in 2012.

Understanding the exact timing of the decline has been important for understanding possible driving forces of this change. The evidence presented in Figure 2.1 suggests that the decline of the

[2]As my aggregate measure of the non farm payroll labor share, I use the most recent data from the NIPAs, which reflects the comprehensive revision of the national accounts that occurred in July 2013.

[3]I begin in 1959, as separate estimates of value added in government enterprises are not available until this time.

Figure 2.2: U.S. Private Nonfarm and Private Nonfarm Business Payroll Labor Shares, 1959-2012



The solid line represents the payroll labor share in the private non-farm sector, the construction of which is described in the text. The dashed line represents the payroll labor share in the private non-farm business sector, which is constructed as is the payroll labor share for the private non-farm sector, but with the exclusion of value added and compensation in non profit institutions serving households.

labor share is almost exclusive to the decade of the 2000s. However, this conclusion is somewhat influenced by the inclusion of non-profit institutions serving households, whose value added is both increasing as a share of total value added and is largely comprised of the compensation of employees. For comparison, Figure 2.2 plots the payroll labor share for both the private nonfarm and private nonfarm business sectors, the latter excluding these non-profit institutions. In contrast to the non-farm sector, there appears to be a slight decline in the payroll labor share for the non-farm business sector, beginning potentially as early as the 1980s. However, even with slight trend decline, the majority of the recent decline is still concentrated in the most recent decade.

2.3 Linking Labor Share Changes to Occupational Tasks

2.3.1 Occupational Labor Shares

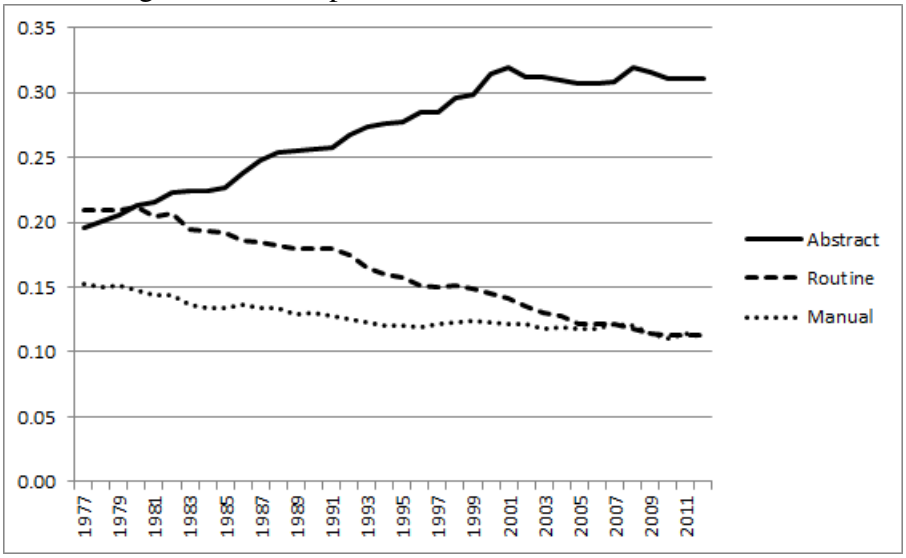
Given the construction of the payroll labor share for the private nonfarm sector, in this section I use household data from the Current Population Survey to decompose aggregate labor compensation by compensation per occupation for occupations specializing in performing different tasks. I decompose the labor share across three groups of occupations - abstract, routine and manual - where each group is defined by the task predominantly performed in those occupations. This occupational grouping follows the definitions of tasks and occupational measures used in Autor and Dorn (2013). Additional explanation of these different types of tasks can be found there. I measure total labor compensation in each occupation group using data on wage/salary income from the March CPS, where consistent occupational codes and the mapping from occupations to tasks is taken from Autor and Dorn (2013). There are a number of challenges in generating total labor compensation using data from the Current Population Survey, including accounting for additional non-wage/salary compensation and top-coding of income data - I document how I account for these in the Appendix.

I decompose the payroll labor share for the non-farm business sector into occupational labor shares for abstract, routine and manual occupations, where the occupational labor share is defined as total compensation for a given occupation divided by total output. This means the labor share can be written as follows:

$$\text{Labor Share} = \frac{(wH)_t}{Y_t} = \frac{(w_a H_a)_t}{Y_t} + \frac{(w_r H_r)_t}{Y_t} + \frac{(w_m H_m)_t}{Y_t} \tag{2.1}$$

Figure 2.3 plots this decomposition for abstract, routine and manual occupations from 1977 to 2012. Despite the fact that the aggregate labor share is roughly constant for much of this period, each of these three shares of income show significant trend behavior throughout the sample. Both the routine and manual labor shares begin to decline in the early 1980s, though the decline for manual occupations is of a lower magnitude and appears to cease midway through the 1990s,

Figure 2.3: Occupational Labor Shares, 1977 to 2012



Occupational labor shares are defined as total compensation paid to workers in a given occupation divided by aggregate output. As such, all three occupational labor shares sum to the total labor share. Data for occupational compensation is obtained from the March CPS and output data is for the private non-farm sector, generated from NIPA data. Additional details regarding construction of occupational labor shares can be found in the Appendix.

with a slight decline again in the Great Recession. Meanwhile, the abstract labor share is steadily increasing until the year 2000, after which it remains nearly constant through the end of the sample. Thus, in terms of this decomposition, the recent decline of the labor share since the year 2000 is primarily about the decline of the routine labor share, as abstract and routine labor shares have been relatively constant over this time period. What is especially noteworthy, however, is that the decline of the routine labor share is not unusual to the last decade - it has been ongoing since the 1980s. Instead, the most notable change in the most recent decade is the sudden slowdown of growth in the abstract labor share. Table 2.1 shows this more concretely by reporting the change in each of these occupational labor shares for 5 year intervals starting in 1977. What we observe is that there is little change in the total labor share through the year 2002, as steady increases in the abstract labor share offset declines in the routine and manual labor shares. But after the year 2002, the routine labor share has continued to decline, while the abstract labor share has stopped increasing, leading to a net decline in the overall labor share.

One interpretation of these facts could be that the the forces driving polarization - and thus,

Table 2.1: 5 Year Percentage Point Changes in Occupational Labor Shares, 1977-2012

	ΔAbs	ΔRout	ΔManual	Δ**Total**
1977-1982	2.7	-0.3	-0.8	**1.6**
1982-1987	2.5	-2.1	-1.0	**-0.6**
1987-1992	2.0	-1.0	-0.8	**0.2**
1992-1997	1.8	-2.5	-0.4	**-1.2**
1997-2002	2.7	-1.5	-0.1	**1.2**
2002-2007	-0.4	-1.3	0.0	**-1.7**
2007-2012	0.2	-0.9	-0.9	**-1.7**
1977-2012	11.5	-9.7	-4.1	**-2.3**
2002-2012	-0.2	-2.2	-0.9	**-3.4**

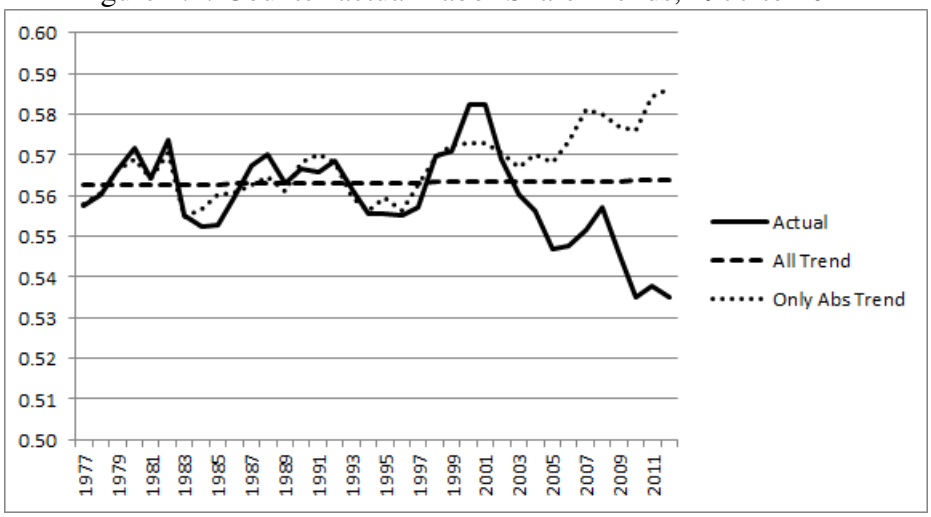
Occupational labor shares are constructed as described in Figure 2.3.

the routine labor share - are primarily responsible for the recent labor share decline, as domestic labor input has been replaced by foreign labor or technological change. This is consistent with the existing theories suggested for the decline of the labor share. Of course, the challenge these theories encounter is that the labor share decline is most pronounced in the 2000s, whereas these theories would predict declines dating back to the 1980s and 1990s, when trade and technology forces began to rapidly increase. What this occupational decomposition highlights is that the reason that this decline only emerges in the 2000s is because it was masked for some period of time by the rising share of abstract income. But since the year 2000, growth in the abstract labor share has substantially slowed while the decline of the routine labor share has continued at a near similar pace. To highlight this point, consider how the labor share would have behaved if each of these occupational labor shares continued their trend behavior observed through the year 1994.[4] Figure 2.4 plots the payroll labor share for the private non-farm sector against two counterfactual labor shares - one where all three occupational labor shares follow a linear trend fit to the period 1977 to 1994 and extrapolated forward to 2012, and one where only the abstract labor share follows a linear trend fit to 1977-1994 and then extrapolated forward to 2012.[5] What we observe is that if all

[4]I choose 1994 to avoid any potentially unusual movements surrounding the 2000s dot com bust, but the results are qualitatively robust to choosing virtually any date prior to 2001.

[5]The qualitative conclusions of this graph are robust to using an exponential trend instead of a linear trend.

Figure 2.4: Counterfactual Labor Share Trends, 1977 to 2012



The solid line presents the payroll labor share in the private nonfarm sector; its construction in Section 2. The dashed line presents the counterfactual payroll share where the occupational labor shares of Figure 2.3 are replaced by linear trends, fit to 1977-1994, and then extrapolated to 2012 and summed together. The dotted line presents the counterfactual payroll share where only the abstract occupational labor share is replaced with such a linear trend, and is added to the data measured routine and manual occupational labor shares.

three labor shares continued their trend behavior from the early part of the sample, the labor share would have truly remained effectively constant for the entire sample. And if only the abstract labor share had followed its trend behavior, we would have observed even a slight increase in the labor share at the end of the period. The point is that without a slowdown of the abstract income share, there is no observed decline of the aggregate labor share. Thus, to understand the recent decline in the labor share, one would have to understand both the ongoing forces driving polarization, as well as the forces that have led to the slowdown in the growth of abstract income.

2.3.2 Controlling for Industry Composition

However, this is only one possible interpretation of these facts. While the decomposition presented above is exact, it can only provide suggestive evidence as to the role of occupational task composition in the labor share decline because it does not control for changes in the task composition of output. Generally speaking, it is not clear how one would measure the task composition of output,

but supposing for the moment that it were possible, then the above decomposition of the labor share could be written as:[6]

$$\text{Labor Share} = \frac{(wH)_t}{Y_t} = \frac{Y_{at}}{Y_t}\frac{(w_aH_a)_t}{Y_{at}} + \frac{Y_{rt}}{Y_t}\frac{(w_rH_r)_t}{Y_{rt}} + \frac{Y_{mt}}{Y_t}\frac{(w_mH_m)_t}{Y_{mt}}$$

With this decomposition, the labor share is related to both the task composition of output as well as the share of total income paid to labor for each task. Thus, the occupational shares of income shown in Figure 2.3 could be driven by either changes in the share of task specific income paid to labor or changes in the fraction of output generated by each task. An alternate interpretation, then, of the decline in the routine labor share is that the share of output generated by routine tasks, $\frac{Y_{rt}}{Y_t}$, has been falling over time, whereas the share of income paid to labor in that task, $\frac{(w_rH_r)_t}{Y_{rt}}$, has been constant. Then, supposing the share of output in abstract tasks was offsetting this decline for the entirety of the sample, the decline in the aggregate labor share would be driven by a decline in the share of abstract income paid to labor. In this case, the decline in the aggregate labor share would have nothing to do with routine occupations. And this is only one of many such possible alternate interpretations.

As measurement of task outputs is not possible (or maybe even sensible), I start by showing that these patterns of occupational labor shares are unchanged when accounting for the industry composition of output. There is considerable variation in the trends and levels of the composition of occupational tasks and the share of total output across industries and thus changes in industry composition can partially proxy for changes in the task composition of output. For a set of I industries, the labor share can be decomposed by industry and occupation as follows:

[6]One approach might be to classify all inputs into production, both labor and non-labor, by the type of task they perform and then decompose income paid to each input by task. But this would present a number of difficulties, including the standard difficulties with measuring capital income, as well as the added task of carefully classifying the task contribution of each type of capital.

$$\text{Labor Share} = \frac{(wH)_t}{Y_t} = \frac{(w_a H_a)_t}{Y_t} + \frac{(w_r H_r)_t}{Y_t} + \frac{(w_m H_m)_t}{Y_t} = \sum_i^I \frac{Y_{it}}{Y_t} \frac{(w_{ia} H_{ia})_t}{Y_{it}} + \sum_i^I \frac{Y_{it}}{Y_t} \frac{(w_{ir} H_{ir})_t}{Y_{it}} + \sum_i^I \frac{Y_{it}}{Y_t} \frac{(w_{im} H_{im})_t}{Y_{it}} \quad (2.2)$$

As accounting for industry output composition requires a separate source of data for industry output, I take data from the GDP by Industry program at the BEA.[7] An empirical challenge of this exercise is having consistent industry codes across the CPS data and the BEA data, however, using the time consistent industry definitions for CPS data of Ruggles et al. (2010) and Census crosswalks between Census industry codes and NAICS industry codes, I am able to obtain a set of 35 time consistent detailed industries and 11 major industries in the non-farm business sector for which I can compute industry specific occupational labor shares.[8] In particular, the 11 major industries I present are very close to the 11 major industries used by Elsby et al. (2013) in their analysis of the labor share, allowing for comparison with the results they obtain at this level of industry disaggregation. For the immediate exercises that follow, I use the 11 major industries for consistency and brevity, though the results presented also hold when using more detailed industry data.

The question is - when industry output shares are held fixed, do we still observe the occupational labor share patterns of Figure 2.3? It is important to observe that this is different from determining what fraction of the change in occupational labor shares is driven by changes in industry composition versus changes within individual industries. We would undoubtedly expect that a nontrivial fraction of the trends in occupational labor shares observed above can be accounted for by changes in industry composition. For example, manufacturing industries are fairly routine

[7] As GDP by Industry data prior to 1998 is yet to be updated to reflect the recent comprehensive revision of the NIPAs, for consistency's sake, I use pre-revision data and scale industries to match the revised data totals. This also constraints the data to only go through 2011. Once the revisions are extended further back, it will be possible to update the tables and figures presented below to represent the most recent data.

[8] Additional details of this industry mapping can be found in the Appendix.

labor intensive and have also been observing a decline in output shares for some time now. But it has been well-documented that changes in industry composition have been largely irrelevant for the most recent changes in the labor share - this is a conclusion reached by Estrada and Valdeolivas (2012), Karabarbounis and Neiman (2014), and Elsby et al. (2013) and holds true for the data used here. Thus, regardless of how much industry composition contributes to changes in the occupational labor shares, these shares will still add up to produce the decline in the labor share even controlling for industry composition. However, there is no guarantee that the trends observed in Figure 2.3 will look remotely the same once industry composition is accounted for. It could be that the meaningful trend differences across occupations that are suggestive of a connection between occupational tasks and the labor share disappear completely when composition is controlled for.

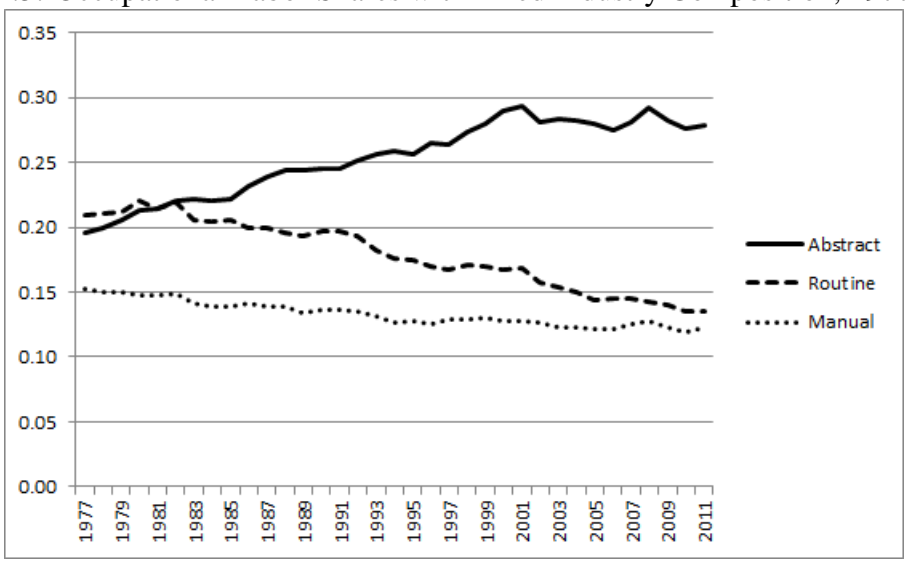
Figure 2.5 presents the occupational labor shares from 1977 to 2011, but with industry output shares held constant at their 1977 level throughout the sample. Even with industry composition accounted for, the same general patterns emerge in the data as those seen in Figure 2.3 - a steady decline in the routine labor share, a slowdown of abstract income in the 2000s, and a relatively flat manual labor share. Similarly, Table 2.2 shows the 5 year changes in occupational labor shares with industry composition fixed and reaches similar conclusions as Table 2.1.[9] Thus, even controlling for industry composition changes, the patterns of occupational income shares are still consistent with interpretation given above.

2.3.3 Occupational Labor Share Changes Within Industries

Although the trends of Figure 2.3 are still visible with only within industry movements, as shown in Figure 2.5, this doesn't necessarily imply that there isn't a change in the task composition of individual industry outputs. A first check against this concern is simply to confirm that sectors contributing substantially to the labor share decline or the recent acceleration of the labor share decline are operating through the routine and abstract task channels, respectively. In this section, I

[9]Notably, there is some difference for the period between 1977 and 1982, a period when industry output share movements were apparently relatively more important.

Figure 2.5: Occupational Labor Shares with Fixed Industry Composition, 1977 - 2011



Construction is similar to that of Figure 2.3, although is based on (2.2), with industry output shares held constant at the year 1977.

Table 2.2: 5 Year Percentage Point Changes in Occupational Labor Shares, Fixed Industry Composition, 1977-2011

	ΔAbs	ΔRout	ΔManual	**ΔTotal**
1977-1982	2.5	1.0	-0.4	**3.1**
1982-1987	1.9	-1.9	-1.0	**-1.0**
1987-1992	1.2	-0.6	-0.3	**0.3**
1992-1997	1.3	-2.6	-0.7	**-2.0**
1997-2002	1.7	-1.1	-0.2	**0.5**
2002-2007	0.0	-1.3	-0.1	**-1.4**
2007-2011	-0.3	-1.4	-0.3	**-1.6**
1977-2011	8.3	-7.5	-3.0	**-2.1**
2002-2011	-0.3	-2.3	-0.4	**-3.0**

Construction of occupational labor shares with fixed industry composition is described in Figure 2.5.

show that this is the case.

Table 2.3 reports changes in occupational labor shares by industry for the 11 major industries over the period 1987-2011, with the period chosen so as to be comparable with the industry level analysis done in Elsby et al. (2013), though the results are robust to alternate periods. The first four columns report the total change in occupational shares and the total labor share by industry, with the fifth column showing the total decline in the labor share multiplied by the industry output share in the year 1987, giving a weighted decline which sums to the aggregate. A few observations from this table are worth mentioning. First, as in Elsby et al. (2013), the industries which contribute the most to the total decline of the labor share are the manufacturing industries and the trade/transportation and utilities industries. Second, both these industries which have seen the greatest labor share decline and the industries which have generally seen substantial labor share declines are those which have experienced the greatest declines in their routine labor shares. The largest decline in the routine labor share is in non-durable manufacturing, which is also has the largest decline in the labor share of any industry. Computing a simple correlation between industry labor share declines and industry routine labor shares, one gets a value of 0.91, suggesting a strong relation between these two; even removing manufacturing industries from the sample, the correlation is still 0.86. Of course, it is possible that this relationship is somewhat of a tautology - wherever labor shares declines are large, by definition, declines in at least some occupational labor share must be large as well. And we do observe that routine labor share has seen the least growth of all occupational labor shares for in almost every sector, suggesting that polarization of income growth is widespread. In principle, however, there could be large routine labor share declines and large abstract labor share increases in a given industry which offset each other which would reduce this measured correlation, and these are not observed.

Additionally, Table 2.4 reports the change in industry occupational labor shares in both the periods 1989-1998 and 2002-2011, as well as the difference between these two.[10] As has already

[10]The two windows are chosen to be of comparable length and excluding the large spike observed surrounding the early 2000s recession. However, using different windows does not substantially change the results.

Table 2.3: Occupational Labor Share Changes by Industry, 1987-2011

	ΔAbs	ΔRout	ΔMan	**ΔTot**	**Wgt. ΔTotal**
Forestry/Mining	-2.3	-5.3	-0.7	**-8.3**	**-0.2**
Construction	6.4	1.2	4.1	**11.7**	**0.7**
Dur. manuf.	5.8	-10.8	-3.7	**-8.8**	**-1.2**
Non-dur. manuf.	-0.2	-20.1	-4.8	**-25.1**	**-2.5**
Trade/Trans./Utilities	1.1	-3.2	-1.3	**-3.4**	**-0.9**
Information	-1.4	-7.7	-3.5	**-12.7**	**-0.5**
Prof/Business services	10.3	-1.9	-1.7	**6.7**	**0.8**
Financial activities	6.0	-6.0	-0.9	**-1.0**	**-0.1**
Ed./Health services	8.5	-1.8	1.9	**8.5**	**0.6**
Leisure/Hospitality	4.3	-1.0	-4.5	**-1.2**	**-0.2**
Other Services	8.4	-0.4	3.3	**11.3**	**0.3**
TOTAL	6.2	-7.2	-2.2	**-3.2**	**-3.0**

The table present changes in total industry compensation in each occupation divided by output of that industry. The final column shows the change in the within industry labor share multiplied by its 1987 output share to highlight the significance of individual sectors' contributions to the aggregate decline in the labor share.

been observed, the majority of the labor decline is occurring in this latter period, and thus if this is directly related to the slowdown of the abstract labor share, we would expect that industries contributing to the difference in labor share changes across these two periods to be observing big difference in their abstract labor share changes. Table 2.4 shows that such a correlation does exist. The key observation to make in Table 2.4 is in the last set of columns, showing the difference in the changes across the two periods. The industries which have contributed the most to the difference in the total labor share change, shown by the weighted total column, are trade/transportation/utilities, professional and business services, and education and health services. Notably, this is a somewhat different set of industries than those driving the total labor share decline, and suggests that understanding the labor share decline may extend beyond understanding changes in manufacturing and trade alone. Further, in each of these industries the driver of this change is the abstract labor share - these industries, along with the information sector, are the industries which witness the largest difference in the change of the abstract labor share. Computing a simple correlation coefficient between the total change in differences and the total change in differences in abstract labor share across industries, one obtains a value of 0.95; removing education and health services, where the

change is largest, the correlation is still 0.89.

In contrast to the results of Table 2.3, there is much more asymmetry across industries in the relative behavior of the abstract labor share growth. In Table 2.3, although there was a clear link between industries with large declines in the routine labor share and industries with large labor share declines, it was obvious that relative to changes in other occupations, the routine labor share was falling in nearly every industry. However, from Table 2.4, we see that there is not a sudden slowdown in abstract labor share growth in every industry, but only in the trade, professional/business, information and education/health sectors - precisely those sectors which contribute the most to the differences in the aggregate labor share change. Thus, there is already evidence here to suggest that the recent change in labor share growth is directly related to this change in the rate of growth of the abstract labor share.

2.3.4 Within Industry Relationships Between Tasks and Labor Shares

The previous section shows that the industries that have observed the largest labor share declines have observed the largest routine labor shares declines, but, as stated above, this evidence has the potential to be circular in nature. Given that polarization seems to be occurring in nearly every industry, a useful question is - how does variation in polarization and/or routine task intensity correlate with the changes in labor shares over time? If indeed, changes in routine and abstract tasks are responsible for the overall decline of the labor share, then we would expect that variation in task intensity and task changes correlates closely with variation in labor share movements at the industry level. Thus, in this section, I consider cross-sectional industry regressions between measures of levels and changes in routine task intensities and changes in labor shares for both the 11 major industries of the previous section and 35 more detailed industries.

Before presenting these results, it is important to make a few observations about the cross-sectional strategy here. First, both Solow (1958) and more recently, Young (2010), observed that there was substantial variation in industry level labor shares in periods when the aggregate labor share remained constant. These earlier results are certainly evident in the results presented in

Table 2.4: Occupational Labor Share Differences in Changes Between 1989-1998 and 2002-2011

	1989-1998					2002-2011					Diff.				
	ΔAbs	ΔRout	ΔMan	**ΔTot**	**Wgt. ΔTot**	ΔAbs	ΔRout	ΔMan	**ΔTot**	**Wgt. ΔTot**	Abs	Rout	Man	**Tot**	**Wgt. Tot**
For/Min.	1.4	1.0	0.5	**2.8**	**0.1**	1.6	-0.9	-1.9	**-1.2**	**0.0**	0.2	-1.9	-2.3	**-4.0**	**-0.1**
Constr.	-1.5	-0.5	-3.5	**-5.5**	**-0.3**	4.8	1.1	2.9	**8.8**	**0.5**	6.4	1.5	6.4	**14.3**	**0.8**
Dur Mfg	4.6	-1.8	-1.1	**1.6**	**0.2**	3.9	-3.0	-1.1	**-0.2**	**0.0**	-0.6	-1.2	0.1	**-1.8**	**-0.2**
N-dur Mfg.	-1.4	-7.4	-0.6	**-9.5**	**-0.9**	-1.2	-7.2	-0.9	**-9.3**	**-0.9**	0.2	0.2	-0.3	**0.2**	**0.0**
Trade etc.	1.6	-1.6	0.0	**0.1**	**0.0**	-2.4	-1.9	-0.7	**-5.1**	**-1.4**	-4.1	-0.4	-0.8	**-5.2**	**-1.4**
Info.	4.5	-3.0	1.8	**3.3**	**0.1**	-4.9	-2.8	-1.3	**-9.0**	**-0.3**	-9.4	0.1	-3.0	**-12.2**	**-0.5**
Prof/Bus.	7.1	-0.9	0.0	**6.1**	**0.7**	-0.8	0.2	-0.5	**-1.2**	**-0.1**	-7.9	1.0	-0.5	**-7.3**	**-0.9**
Finance	1.7	-3.1	-0.7	**-2.1**	**-0.3**	2.7	-1.7	0.0	**1.0**	**0.1**	1.0	1.4	0.8	**3.2**	**0.4**
Ed./Health.	16.1	-0.2	1.6	**17.5**	**1.3**	-9.2	-1.9	0.1	**-10.9**	**-0.8**	-25.3	-1.7	-1.4	**-28.4**	**-2.1**
Leis/Hosp.	-0.3	-2.4	-1.5	**-4.2**	**-0.1**	0.8	-0.2	-2.1	**-1.5**	**0.0**	1.1	2.2	-0.6	**2.7**	**0.0**
Oth. Serv.	1.7	-1.9	-3.6	**-3.9**	**-0.1**	1.0	-0.4	2.8	**3.3**	**0.1**	-0.6	1.5	6.4	**7.2**	**0.2**
TOTAL	4.1	-2.8	-0.6	**0.7**	**0.7**	-0.2	-2.2	-0.7	**-3.1**	**-2.9**	-4.3	0.6	-0.1	**-3.8**	**-3.6**

The table present changes in total industry compensation in each occupation divided by output of that industry for the two periods 1989-1998 and 2002-2011, as well as the differences in the changes between these two periods. The weighted columns show the change in the within industry labor share multiplied by its 1987 output share to highlight the significance of individual sectors' contributions to the aggregate decline in the labor share.

Table 2.4, where there have been significant changes in industry labor shares even in the early period of 1989-1998 when the aggregate labor share remained virtually constant. As such, even if the aggregate labor share changes are purely about changes in factor shares across abstract, routine and manual tasks, there is likely to be a lot of cross industry variation that will be unexplainable by these measures. Second, given that there is quite a bit of variation in labor shares across industries, if the underlying drivers of aggregate changes are limited to just a few sectors, it will be quite difficult to identify these off of cross-sectional variation. This is unlikely to be a problem for relating long run declines in labor shares and changes in routine tasks, as the routinization process has been observed in nearly every industry, something Table 2.3 corroborates.[11] However, as observed in Table 2.4 and discussed above, it does not appear that changes in abstract tasks have been nearly so uniform. This suggests that using cross industry variation may not be informative for understanding those changes. This is particularly true for using more detailed industries, as the added detail mostly occurs in manufacturing, which in the aggregate, has shown far less of a contribution of recent changes in abstract and total labor declines. Thus, in this section, I only focus on the relationship between long run labor share declines and variation in polarization or routine task intensity.

I consider three measures of routine task intensity - the share of total labor income paid to routine workers, the share of hours worked in routine occupations, and the routine wage relative to the average industry wage - and I use both the initial levels of these measures and their changes from 1987-2011.[12] Tables 2.5 (11 industries) and 2.6 (35 industries) report regressions of changes in industry labor shares on these measures using a time period of 1987-2011, again, for comparability with Elsby et al. (2013).[13] Between these six measures of routine task levels and changes, there is a substantial amount of explanatory power. In columns (1) - (6) of both tables 2.5 and 2.6, we see that every single one of these routine task measures, except for the change in the routine hours

[11]The Appendix of Chapter 1 also shows this point.

[12]I have tried alternate measures of relative wages and found that the conclusions do not vary significantly with these choices.

[13]And again, the results are not particularly sensitive to this choice of timing.

share in the 11 sector regressions, is significant at the 10% level, and many are significant at the 5% or 1% levels, particularly the change in routine labor income shares and initial routine income and hours shares. As these three routine task measures show the strongest relationship with the labor share declines, I will focus my attention on these. In addition to their statistical significance, there is a substantial amount of variation explained by these single measures in each regression - for changes in routine labor income shares and initial routine income and hours shares, the regression R^2s are 0.52, 0.72, 0.78 and 0.18, 0.23 and 0.24 for the 11 sector and 35 sector regressions respectively.

One possible concern, particularly for the initial shares, is that these routine task measures are simply a proxy for the manufacturing industry, where there have been substantial labor share declines. To account for this, the regressions reported in columns (7)-(12) include a dummy variable for being a manufacturing industry. I observe that for the change in routine income share and initial routine income and hours shares, the inclusion of the manufacturing indicator does not undo the statistical significance of the routine task measure, and in fact, the manufacturing indicator is actually insignificant for these specifications. Thus, although manufacturing industries tend to be more routine intensive, the connection between labor share changes and routine tasks goes beyond this sector.

Finally, columns (13)-(15) in Tables 2.5 and 2.6 report the combination of the change and level variables for each of the three routine task measures. Again, we see statistical significance for many of these regressors and substantial explanatory power in these regressions for explaining variation in industry labor share changes. The conclusion reached from these exercises is that there is substantial evidence that industries which have experienced greater polarization, as measured by changing labor income shares, and industries which were initially more routine intensive, have experienced greater labor share declines over the past 25 years. Thus, there is strong evidence consistent with the hypothesis that the central driving force behind the long run decline of the labor share is the ongoing process of routinization and polarization.

Of course, identifying these patterns from a cross-sectional regression of industries doesn't

Table 2.5: Regressions of Industry Labor Share Changes on Industry Routine Intensity Measures, 1987-2011 (11 Industries)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
ΔRout. Inc. Share, '87-'11	1.32**						0.854*						0.612**		
	(0.533)						(0.443)						(0.265)		
ΔRout. Hr. Share, '87-'11		1.08						0.716						0.247	
		(0.848)						(0.481)						(0.466)	
$\Delta\frac{\text{Routine Wage}}{\text{Average Wage}}$, '87-'11			0.783*						0.493*						1.051***
			(0.411)						(0.246)						(0.249)
Rout. Inc. Share, '87				-0.709***						-0.665***			-0.560***		
				(0.140)						(0.124)			(0.114)		
Rout. Hr. Share, '87					-0.618***						-0.543***			-0.583***	
					(0.139)						(0.104)			(0.139)	
$\frac{\text{Routine Wage}}{\text{Average Wage}}$, '87						-0.502*						-0.341			-0.798***
						(0.263)						(0.171)			(0.197)
Mfg Indicator							-0.112**	-0.158**	-0.139*	-0.018	-0.033	-0.162**			
							(0.050)	(0.050)	(0.064)	(0.065)	(0.063)	(0.067)			
R^2	0.52	0.22	0.37	0.78	0.72	0.18	0.66	0.60	0.64	0.78	0.73	0.59	0.85	0.73	0.80

Significance levels: * p<10%, ** p<5%, *** p<1%. All regressions are weighted using average value added share for 1987-2011, though results are qualitatively robust under unweighted regression. Results are also robust to a variety of methods of accounting for potential outliers. All shares are measured as fractions between 0 and 1.

Table 2.6: Regressions of Labor Share Changes on Routine Intensity Measures, 1987-2011 (35 Industries)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)
ΔRout. Inc. Share, '87-'11	0.975***						0.719***						0.555**		
	(0.268)						(0.242)						(0.257)		
ΔRout. Hr. Share, '87-'11		0.747*						0.585*						0.346	
		(0.372)						(0.341)						(0.359)	
$\Delta\frac{\text{Routine Wage}}{\text{Average Wage}}$, '87-'11			0.337*						0.219						0.282
			(0.222)						(0.194)						(0.218)
Rout. Inc. Share, '87				-0.632***						-0.417**			-0.381**		
				(0.219)						(0.196)			(0.146)		
Rout. Hr. Share, '87					-0.467***						-0.383**			-0.424***	
					(0.129)						(0.183)			(0.133)	
$\frac{\text{Routine Wage}}{\text{Average Wage}}$, '87						-0.287*						-0.199			-0.217
						(0.264)						(0.255)			(0.276)
Mfg Indicator							-0.097	-0.131**	-0.131**	-0.047	-0.048	-0.136**			
							(0.063)	(0.056)	(0.061)	(0.086)	(0.083)	(0.060)			
R^2	0.18	0.08	0.06	0.23	0.24	0.05	0.24	0.21	0.18	0.25	0.25	0.18	0.29	0.25	0.09

Significance levels: * p<10%, ** p<5%, *** p<1%. All regressions are weighted using average value added share for 1987-2011, though results are qualitatively robust under unweighted regression. Results are also robust to a variety of methods of accounting for potential outliers. All shares are measured as fractions between 0 and 1.

give a clear implication for what fraction of the aggregate labor share decline is attributable to routinization and polarization, and this hypothesis is naturally subject to the timing concerns discussed in the introduction. The key question that remains, then, is explaining why there has been such a slow down in the growth of the abstract labor share. As already mentioned, there are a number of reasons why cross-sectional regressions are less likely to be effective in confirming the relationship of changes in abstract tasks and changes in effectively the growth rate of the labor share.[14] This suggests that a more pertinent immediate focus would be on those specific industries where there is this notable slowdown of abstract labor share growth, as identified in the prior section, and using more detailed data to ascertain the source of this change.

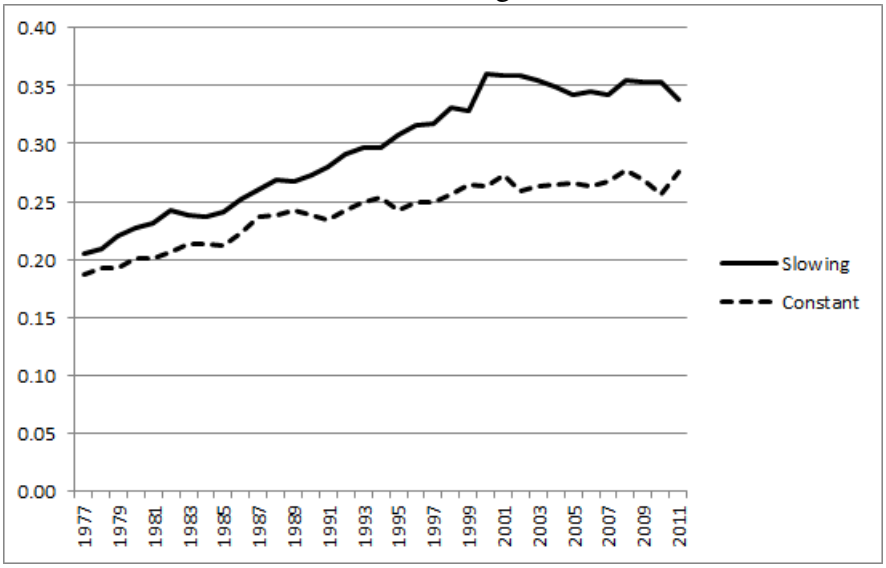
2.4 Understanding the Slowdown of Abstract Income

2.4.1 Occupational Differences Between Sectors

Whereas labor market polarization has been occurring in almost every industry, the recent slowdown of growth in abstract income appears to be limited to a specific subset of industries, as seen previously in Table 2.4, occurring in only four of the eleven major industries - trade/transportation/utilities, professional/business services, information services, and education/health services. For what follows, I will combine these four sectors into one unified sector and call it the slowing sector, and compare this sector of the economy with the residual sectors (natural resources/mining, manufacturing, construction, financial activities, leisure/accommodation, other services), which I will refer to as the constant sector. Figure 2.6 reports the abstract labor share in these two sectors from 1977-2011 and highlights the contrast between the two. The abstract labor share in the slowing sector is indeed slowing - between 1989 and 1998, its abstract labor share increased by about 6 percentage points, but from 2002 to 2011, the labor share declined by 2 percentage points. In contrast, between 1989 and 1998, the constant sector labor share increased by 1.4 percentage points and

[14] A similar series of regressions using indicators of abstract tasks and growth in abstract tasks to explain changes in the rate of change of the labor share finds limited evidence of a consistent pattern between the two - detailed results are available from the author by request.

Figure 2.6: Abstract Labor Share in Slowing and Constant Sectors, 1977-2011



The abstract labor share is given by total compensation paid to abstract occupations divided by total output for a given sector. The slowing sector is made up of: trade/transportation/utilities, professional/business services, information services, and education/health services. The constant sector is made up of natural resources/mining, construction, durable and non-durable manufacturing, financial services, leisure/accommodation services, and other services.

between 2002 and 2011, it increased by 1.6 percentage points.[15]

Why has growth in the abstract labor share slowed so much in the slowing sector? Additional insight can be gained by decomposing this slowdown into detailed abstract occupations. I consider a division of abstract occupations into ten occupational groups: managers and management related occupations, engineers, computer scientists and programmers, scientists, medical occupations (doctors/nurses/therapists), teachers, lawyers/social workers/social scientists, sales occupations, technicians, and arts/media/entertainment occupations.[16] A detailed description of how these are derived from the consistent occupational codes of Autor and Dorn (2013) is given in the Appendix. I decompose the abstract labor share into detailed abstract occupation labor shares in the same manner as the aggregate labor share was decomposed originally in (2.2). Because some of

[15]These results are robust to holding fixed the underlying detailed industry composition.

[16]The sales occupations included in abstract labor are a specific subset of sales occupations - financial, insurance, real estate sales, as well as sales supervisors or proprietors. As FIRE is in the constant sector, over 95% of individuals in sales occupations in the slowing sector are sales supervisors or proprietors.

the individual occupations show a great deal of fluctuations and the emphasis here is more on long run changes, I smooth each occupation series with an HP filter with parameter 6.25 to avoid added sensitivity to the choice of the time windows. Further, while the following results are reported using output for the whole slowing sector, they are robust to controlling for industry composition.

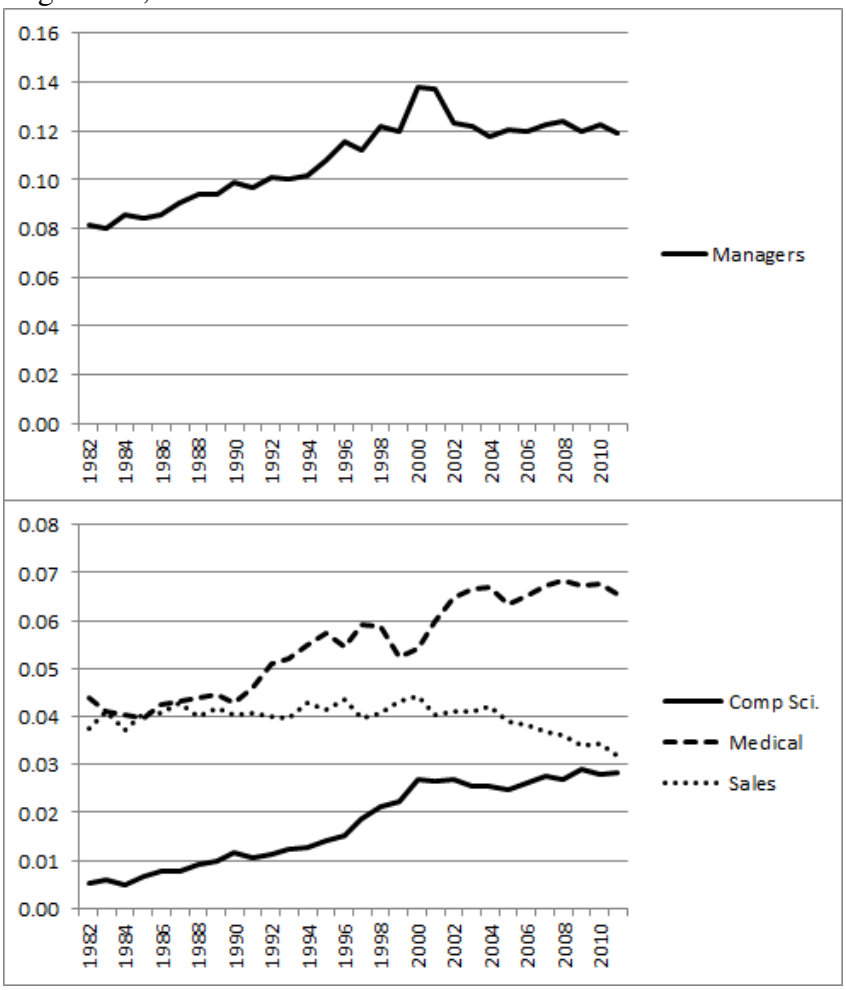
Table 2.7 reports changes in detailed abstract occupation labor shares for the same two periods as in Table 2.4, 1989-1998 and 2002-2011, and the difference between these two changes and each occupation's contribution to the overall change. Although the rate of change in each abstract occupation's labor share has declined some in the most recent period, four occupations contribute disproportionately to the decline - managers, computer scientists/programmers, medical occupations, and sales occupations. Combined, these four occupations account for 86% of the change, with managers alone contributing over half of that at 47%. This suggests that there is significant asymmetry across abstract occupations, and that understanding the abstract slowdown relates to this small handful of occupations.

The specific abstract labor share for each of these three occupations are plotted in Figure 2.7.[17] For managerial and computer science/programming, the labor share in the slowing sector was growing steadily and then rapidly accelerated in the middle of the 90s, before a substantial slowdown in the 2000s recession, after which growth hasn't fully returned to even its pre-1990s rate. Medical occupations show a similar acceleration in the 1990s, although growth in this labor share seemed to begin several years earlier; like managers and computer scientists/programmers, this growth slowed substantially the 2000s. In contrast, the labor share for sales occupations in the slowing sector was virtually constant until the early 2000s, when it began a significant trend decline.

How do these results compare to what has been observed in the constant sector? Table 2.8 shows the comparison between the difference in labor share changes between 1989-1998 and 2002-

[17]The plots in Figure 2.7 begin in 1982 instead of 1977, as there is a not insignificant change in the classification of managers and sales occupation workers which results in a discontinuity between 1981 and 1982. However, aside from this, there are no significant discontinuities within abstract occupational measurement for the remainder of the period.

Figure 2.7: Labor Shares in Abstract Managerial, Computer Science, Medical and Sales Occupations in the Slowing Sector, 1982-2011



Detailed abstract occupation labor shares for the slowing sector are given by total compensation paid to those occupations divided by total output in the slowing sector. Detailed occupations are derived using the consistent occupational codes of Autor and Dorn (2013); income measurement details and the occupation codes comprising each detailed abstract occupation group are reported in the Appendix. Managerial occupations are presented in the upper panel, while computer science/programming, medical and sales occupations are reported in the lower panel.

Table 2.7: Detailed Decomposition of Abstract Labor Share Changes for 1989-1998 and 2002-2011, Slowing Sector

	1989-1998	2002-2011	Change	% Tot. Decline
Managers	2.7	-0.7	-3.4	47.3%
Engineers	0.0	-0.3	-0.3	3.7%
Comp. Sci/Prog.	1.1	0.3	-0.8	12.3%
Scientists	0.0	0.0	0.0	0.2%
Medical	1.3	0.4	-0.8	11.7%
Teachers	0.3	0.2	-0.1	1.9%
Law/Social	0.4	0.2	-0.2	1.9%
Sales	0.1	-0.9	-1.0	14.6%
Technicians	0.0	-0.1	-0.1	1.6%
Other	0.3	-0.1	-0.4	4.9%

Detailed abstract occupation labor shares for the slowing sector are given by total compensation paid to those occupations divided by total output in the slowing sector. Detailed occupations are derived using the consistent occupational codes of Autor and Dorn (2013); income measurement details and the occupation codes comprising each detailed abstract occupation group are reported in the Appendix. The first two columns report the change in the smoothed labor share (HP filter, 6.25) for a given occupation between the periods 1989-1998 and then for 2002-2011, with the third column reporting the difference between these changes.

2011 for each detailed occupation in both the slowing and constant sectors.[18] What we observe is that these four occupations contributing to most of the slowdown in the slowing sector are sectors where there are substantial discrepancies in slowdown between the slowing and constant sectors. In particular, in managerial, computer programming and medical occupations in the constant sector, there is virtually no difference in labor share changes between 1989-1998 and 2002-2011. This is not particularly surprising for medical occupations, as these are primarily concentrated in medical services industries, which are exclusively in the slowing sector. But for computer programming and managerial occupations, this discrepancy is notable as both are substantial components of abstract labor income in each sector. Examining sales occupations, notably, there is a substantial

[18]Notably, summing up the changes in each occupation for the constant sector, there is a nontrivial measured slowdown in the smoothed total labor share in the slowing sector (though still far less than in the slowing sector). This is largely attributable to the strong response of the filtering procedure to the Great Recession, which particularly affected the constant sector, as it includes both the financial and construction industries. However, it is clear from the table that once sales occupations are accounted for (largely financial, insurance and real estate sales), there is no measured slowdown in the smoothed total abstract labor share for the constant sector.

Table 2.8: Difference in Labor Share Changes Between 1989-1998 and 2002-2011, Slowing and Constant Sectors

	Slowing	Constant	Difference
Managers	-3.4	-0.2	-3.2
Engineers	-0.3	0.5	-0.8
Comp. Sci/Prog.	-0.8	-0.1	-0.7
Scientists	0.0	0.1	-0.1
Medical	-0.8	0.1	-0.9
Teachers	-0.1	-0.1	0.0
Law/Social	-0.2	-0.2	0.0
Sales	-1.0	-1.6	0.6
Technicians	-0.1	0.2	-0.3
Other	-0.4	-0.1	-0.3

Detailed abstract occupation labor shares for the slowing sector are given by total compensation paid to those occupations divided by total output in the slowing sector. Detailed occupations are derived using the consistent occupational codes of Autor and Dorn (2013); income measurement details and the occupation codes comprising each detailed abstract occupation group are reported in the Appendix. The first two columns report the difference in the change in the smoothed labor share (HP filter, 6.25) for a given occupation between the periods 1989-1998 and 2002-2011, with the third column reporting the difference between these changes across the two sectors.

decline in the sales occupation labor share in the constant sector, and of an even greater magnitude than in the slowing sector, likely augmented by the Great Recession's impact on financial and real estate sales, which compose a significant fraction of sales in the constant sector. These results further highlight that the decline in the abstract labor share appears to be particularly specific to changes in these particular occupations, and primarily in the slowing sector.[19]

2.4.2 Hours vs. Wages in Abstract Occupational Labor Shares

Before suggesting any possible theories for why there has been a slowdown in income paid to these occupations in these sectors, I present one last piece of evidence regarding these specific changes. I decompose growth in abstract labor shares in the following way - for an occupation j within a given sector i, the occupational labor share can be written as the product of the ratio of the wage in that occupation and industry, w_{ij}, to industry-wide output per hour, y_i/h_i and the share of total

[19]Notably, there is also a significant discrepancy between engineering occupations between the two sectors, which is not easily explained.

sector hours worked in that occupation, h_{ij}/h_i:

$$\frac{w_{ij}h_{ij}}{y_i} = \frac{w_{ij}}{\frac{y_i}{h_i}}\frac{h_{ij}}{h_i} \tag{2.3}$$

Thus, growth in the occupational labor shares can be considered as coming from growth in wages and growth in hours, both relative to industry-wide growth.[20] Table 2.9 presents the growth in the wage and hours component for the slowing and the constant sectors' total abstract labor shares in the periods 1989-1998 and 2002-2011, as well as for major contributing occupations in the slowing sector and reports the contributions of both hours and wages to the total difference in growth across periods. Several observations stand out from this table. First, the slowing sector is seeing both a decline in hours and in wages growth, with three quarters of the total decline occurring in wage growth. For the individual occupations within that sector, all are seeing lower wage growth (negative, even) in the latter period, and most are seeing slower hours growth, with the exception of medical occupations, where hours have continued to grow. Second, comparing to the constant sector, this sector has also seen a substantial slowdown in wage growth (though only just over half as much), but has not seen a corresponding slowdown in hours growth. Thus, one of the stark differences between the two sectors is the difference in the hours growth, even though this only contributes 26% of the total growth slowdown for the slowing sector.

2.4.3 Possible Explanations

In this section, I discuss two possible explanations for this phenomenon - changes in the composition of compensation and changes in the supply and demand for certain types of higher skills.

One possible explanation for these trends pertains to the substantial reversal of wage growth

[20] Since many of these labor shares have a fairly linear trend, looking at log changes will naturally overstate the recent decline. An alternate approach here would be to do counterfactual exercises where either the wage/productivity ratio or the hours share is fixed while the other is allowed to vary and decompose changes in this manner. Using this alternate approach generates almost identical quantitative results regarding the relative importance of wages and hours for the labor share slowdown, and thus I present the growth decomposition, as this is more straightforward.

Table 2.9: Growth in Hours and Wage Components of Occupational Labor Shares, 1989-1998 and 2002-2011

	Hours Growth		Wage Growth		Hours/Wage %
	1989-1998	2002-2011	1989-1998	2002-2011	
Constant Sector	11%	12%	0%	-15%	-7%/107%
Slowing Sector	14%	5%	11%	-12%	26%/74%
Managers	16%	1%	12%	-11%	41%/59%
Comp Sci.	64%	17%	16%	-13%	62%/38%
Medical	12%	22%	16%	-19%	-35%/135%
Sales	0%	-13%	6%	-17%	38%/62%

Growth in hours and wage components for each period are given by log differences of the appropriate component (shown in (2.3)) between the end and start dates of the period. The last column documents the growth differences of the hours and wage components divided by the difference in total growth, where total growth is growth in the hours component plus growth in the wage component. For the results above, the scaling factor for non-wage compensation, as discussed in the Appendix, is omitted, however, results are robust to inclusion. Definitions of sectors and occupations are given in the text and in the Appendix.

in the 2000s. As shown above, wage levels in abstract occupations, relative to industry productivity, were growing or constant up until the 2000s, after which they declined substantially. One possible explanation is that this reversal may simply be a consequence of mismeasurement. In the past several decades, there has been a substantial change in the composition of compensation of employees, especially top executives, where compensation has been paid in ownership rights to the firm, either through stock grants or stock options. This type of compensation is not captured in the Current Population Survey and not well captured in the national accounts either.[21] Thus, if there was a significant change in the composition of compensation towards stocks and other unobserved or unreported forms of income around the year 2000, wages in the data would be understated and as a result, there would appear to be a slowdown in the abstract labor share amongst those occupations receiving compensation in the form of stocks, amongst whom managers and executives are most prominent.

Unfortunately, there is only very limited data on the composition of compensation between

[21]Moylan (2008) shows that stock options are only observed in the compensation data at the time of exercise and only when it is non-qualified stock options. However, even this assumes perfect reporting on the behalf of firms, and there is ample reason - some mentioned in Moylan (2008) - as to why this might be understated in reporting.

cash payments in wages and salaries and stock payments through grants or options.[22] The best available data is for the top paid managers in publicly listed firms, available through ExecuComp. While this data is certainly not representative of all managers or even all occupations, we do see that managerial occupations contribute the most to the slowdown of the abstract labor share and this can be used at least as an initial benchmark to indicate trends in compensation composition. Using this data, I decompose executive compensation into five primary components: salaries, bonuses, long-term incentive payments, the value of stock options granted, and restricted stock grants.[23] This is plotted in Figure 2.8 for the years 1992-2012.

Generally, the trend of total executive compensation appears to follow closely the patterns in managerial compensation for the slowing sector shown in Figure 2.7. But very little of total compensation is paid in salary, which is the primary component of what we see in the Current Population Survey, which is somewhat surprising.[24] This suggests that either the executive series isn't representative of managers at large, or that perhaps some non-wage/salary income is being reported in the CPS.[25] And although the share of salary in total managerial compensation is generally falling throughout this period, this is mostly occurring in the period prior to 2000. Following the

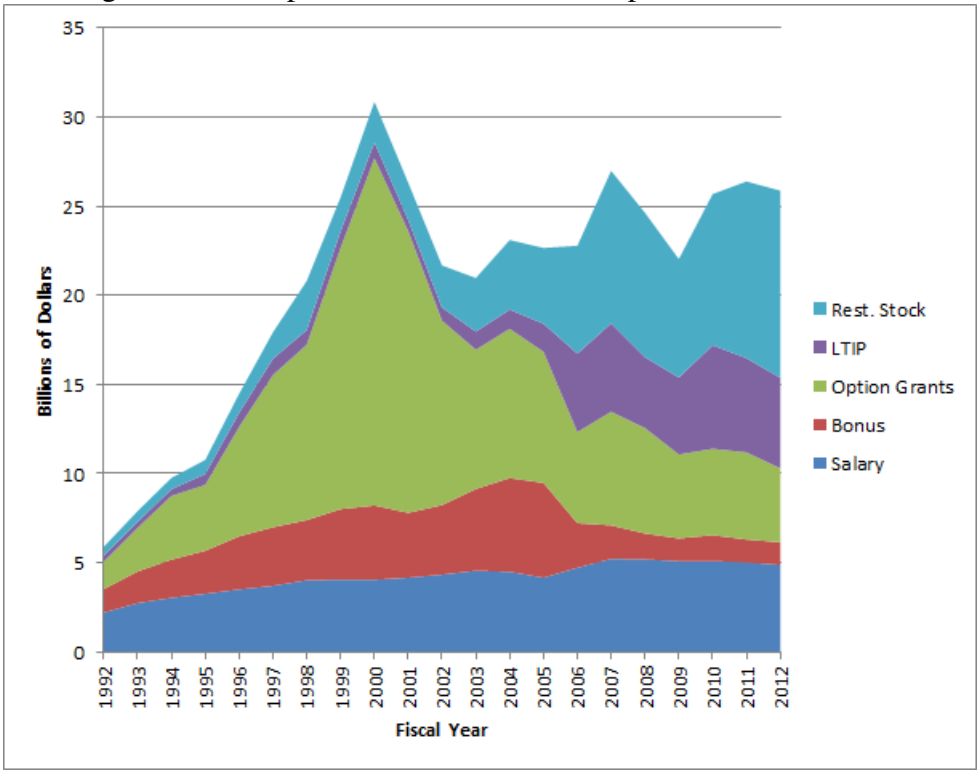
[22]Notably, tax record data, such as that used in Atkinson, Piketty and Saez (2011), is also limited in its ability to decompose compensation. Similar to the national accounts, individual tax records only report income from nonqualified stock options when the options are exercised, not when they are granted. Further, other types of stock compensation, which may differ in their tax treatment, are not easily separable from capital gains income at large. Thus, tax data can perhaps given some additional insight, but it too is limited in being able to speak precisely to this point.

[23]There are some issues with the ExecuComp data around the year 2006; I follow the procedures of Mishel and Sabadish (2013) for constructing a time-consistent series. Even with a time consistent series, there does seem to be some reclassification of bonus payments as long term incentive payments around 2006, and the valuing method for stock options switches from Black-Scholes to fair value at 2006.

[24]Further, some components of both salaries and bonuses can be paid through non-cash means, which may not be accurately reported in the household survey data.

[25]Lambert and Larcker (2001) reported that most employees didn't understand "the basic economics of stock options" and cite studies showing that as much as 52% of employees with stock options knew nothing about their tax implications. If the degree of ignorance is that high, then it's possible that some of this is being reported in the CPS, even though it's not asked for. Also, since for tax reporting, exercising non-qualified stock options appears as wage and salary income, it's possible that respondents using their tax records to respond to CPS questions on income may have included this in their response.

Figure 2.8: Composition of Executive Compensation, 1992-2012



Data is from the ExecuComp database. Executive compensation is defined as the sum of five components: salary, bonus, the value of stock option grants, long term incentive payouts, and restricted stock grants. Construction of these variables from the ExecuComp data follows the procedures of Mishel and Sabadish (2013).

year 2000, the share of compensation in stock options granted actually drops dramatically, though this is partially offset by an increase in compensation through restricted stock grants. These trends are similar to those found in Frydman and Jenter (2010) for CEOs alone. If there is a significant change in compensation composition occurring after 2000, it is coming through a shift away from stock options toward restricted stock grants, and based on differential tax treatment of the granting, exercising and sale of these forms of stock, it is yet possible that this could be important for measured compensation, especially if workers were reporting gains from stock option exercise in the CPS prior to 2000. But it is not clear that a careful exercise separating these two from the underlying data in the national accounts or the Current Population Survey (again, which really shouldn't be picking up any stock compensation) is possible given data restrictions. Thus, at least on the surface, it does not appear that there is strong evidence to suggest that there is a major change in the composition of compensation away from salary in the 2000s.

An alternate explanation for the slowdown of the abstract labor share pertains to changes in the supply and demand for certain skills. Beaudry et al. (2013) have suggested that there was a great reversal in the demand for higher skilled labor - particularly, those employed in abstract occupations - around the turn of the century and this has driven a decline in the growth of employment and wages for these higher skilled workers in the past decade. They argue that this could be driven by the IT revolution reaching a maturity point, which has slowed the need to accumulate higher skills. A theory such as this certainly is consistent with some of the observed facts above: the timing is generally correct and we would expect (some) managers and computer scientists/programmers to be particularly engaged with information technology. But of course, the slowdown in the abstract labor share is limited to certain sectors of the economy, and it is not clear why demand for skills would shift so asymmetrically across industries if information technology is indeed a general purpose technology. Further, it is not as obvious how a slowing in the rapid growth of IT progress would so directly affect doctors or sales occupations. If anything, one might suspect that continued development of IT technology - in particular, the rise of e-commerce in the last decade - could be a reason for reduced demand for certain high skilled sales occupations (supervisors and proprietors)

in the trade industries. That said, a further understanding of how changes in technological progress have affected the demand for higher skill levels across occupations and industries seems like an important avenue to pursue for understanding these changes.

2.5 Conclusion

In this chapter, I have viewed the recent decline of the labor share in the U.S. through the lens of changes in compensation across occupational tasks and reached two broad conclusions. First, there appears to be a robust connection between the labor share decline and the ongoing polarization of the labor market. Second, the forces driving polarization and reducing the share of output paid as compensation to routine workers were offset by an increasing share of output paid to abstract occupations until the year 2000, but not afterward. In particular, the abstract labor share has effectively plateaued since the year 2000, and this sudden change has led to the substantial declines in the aggregate labor share as the decline of routine compensation has continued through the recent decade. Early analysis of this slowdown in the abstract labor share traces the change to managerial, computer programming, medical and sales occupations in the trade, information, professional and business service and education/health services industries. The exact cause of this slowdown in these occupations and industries remains unclear, however. There is limited evidence for explanations based in changes in the composition of compensation, and explanations based on a slowing demand for high skills because of maturity in the IT revolution at best need further refinement to explain why the slowdown has been so asymmetric. Thus, to understand the recent decline of the labor share in the US, further attention should be given to understanding what the drivers of the abstract compensation slowdown are and why they operate asymmetrically across industries and occupations.

2.6 Appendix

2.6.1 Decomposing Total Compensation across Occupations

There are a number of potential challenges in generating an aggregate compensation series comparable to those published by the BLS using purely CPS data. In this section, I document how I work with the CPS data to generate such a series.

Basic data files come from the IPUMS database (King et al. (2010)) and cover the Annual Social and Economic Supplement (March) files from the Current Population Survey for the years 1976-2013, providing a sample for 1975-2012, as each sample reports on the prior year's income. I focus on payroll employees in the private non-farm sector, so I restrict my sample to only individuals who were employed in wage and salary jobs in the prior year and who were employed in non-farm, non-government, and non-private household industries.[26] Because data on whether an individual was employed in a nonprofit or profit industry within the private sector is only available after 1994, it is not possible to accurately extract compensation of non-profit institutions serving households, and thus not possible to further restrict the sample to consider the non-farm business sector.[27]

Labor compensation is measured based on individuals' responses to inquiries regarding their wage and salary income. There are two primary shortcomings of this measure. First, wage and salary income in the CPS is top-coded, which, if unaccounted for, will have a significant impact on the measured total income. To account for this, I use the cell means procedure of Larrimore et al. (2008) and replace the top-coded values in each period prior to 1996 with the average income for individuals above the top code from the internal CPS data, which is the procedure imposed in the data after 1996. In addition to this, I impute a Pareto distribution to the tail of the income distribution to correct for the internal top coding of the CPS. Because data on secondary wage and

[26] If a self-employed individual reports any wage/salary income, this income is not excluded, however.

[27] Even after 1994 when the profit/non-profit division is available in the data, it is still not clear which non-profits are serving households and which are serving business. The former are excluded from non-farm business; the latter are included.

salary earnings is only available after 1987, and because there has been wide variance in the public and internal top codes for this source of income, I apply the cell means procedure to total wage and salary earnings from 1976 to 1995. Additionally, the income distribution for secondary earnings is extremely skewed, which makes careful Pareto imputation difficult, so I use the parameters of the Pareto distribution for primary wage and salary earnings for the secondary earnings imputation. While these procedures allow for a more sophisticated accounting for top coding, using a simpler procedure, such as multiplying income at the public top code by 1.5 in all years, as in Autor et al. (2008), generates no meaningful differences in trends in compensation across occupations.

The second primary shortcoming of measuring labor compensation using wage and salary income data is that in the national accounts, roughly 15% of total compensation comes from supplements to wage and salary income, which is not well measured in the CPS.[28] In terms of wage and salary income, as noted by Heathcote et al. (2010) and McCully (2013), the CPS totals, trends and fluctuations are remarkably close to those in the national accounts, so the primary concern is accounting for non-wage and salary compensation. I proceed under the assumption that the trends in wage and salary compensation across occupations are the same as the trends in non-wage and salary compensation across occupations, and scale wage and salary compensation to match the aggregate compensation for the private non-farm sector. This assumption is not entirely unjustified - starting in 1991, the CPS also collected information regarding employer contributions to health care, one of the largest components of non-wage and salary compensation. Using this shortened time window, one can observe that the patterns in this aspect of compensation closely match the patterns in wage and salary income. Further, using the aggregate growth rate to backcast a series of employer contributions to health care prior to 1991 and adding the combined series to the wage and salary has virtually no effect on the results presented. Additional details regarding this data and these comparisons are available from the author upon request.

For measuring income by occupation, I use the time-consistent occupational codes of Autor and Dorn (2013), as well as their mapping of occupations to tasks to identify abstract, routine

[28] Supplements to wages and salaries are comprised of employer contributions to insurance and pension plans and employer contributions to government social insurance.

and manual occupations.[29] In spite of the great usefulness of the time-consistent occupational codes, there still exist small discontinuities in abstract, routine and manual occupations at the periods when the occupational classification system changed. As such, I smooth these changes by imposing the assumption that the share of income (hours) within a given industry in each of these three occupations is the same across this break by multiplying through by a scaling factor, after which I re-scale total industry income so that it remains unchanged.

Finally, the CPS has been subject to a number of redesigns, changes in population weights and trends in imputed values. A variety of alternate methods of accounting for these has shown to have no meaningful effect on the results presented.

2.6.2 Consistent Industry Codes and Detailed Abstract Occupation Codes

For industry comparisons using income data from the CPS and output data from the BEA, it is necessary to build a consistent set of industries. I begin by using the time-consistent industry codes based on the 1990 Census Occupational Codes available through IPUMS (King et al. (2010)) for the CPS income data.[30] On the output side, I use the 1997 NAICS based time consistent industry measures of value added constructed by the BEA through its GDP by Industry program.[31] Finally, using crosswalks provided online by the Census Bureau, I generate the finest possible set of industries which allow for a consistent mapping of 1997 NAICS codes to 1990 Census Occupational Codes. In total, this is 35 industries, 15 manufacturing and 20 non-manufacturing. The list of these industries, the corresponding 1997 NAICS codes and 1990 Census Occupational Codes are given in Table 2.10.

[29] The codes in Autor and Dorn (2013) do not extend to provide time-consistent coverage for the most recent Census occupational code revision, implemented in the CPS in 2010. As such, I use crosswalks provided by the Census and generate a consistent mapping for the last several years. Details of this precise mapping are available upon request.

[30] Technically, IPUMS does not construct this consistent industry mapping for last year's industry, so this mapping has to be obtained directly using translation tables available from the IPUMS-USA website

[31] Again, these are only available for the pre-revision data and totals need to be scaled to match the post revision data.

Table 2.10: Consistent Industry Mapping Across 1990 Census and 1997 NAICS codes

Major Industry	Detailed Industry	1997 NAICS	1990 Census Codes
Natural Resources/Mining	Forestry, Fishing, Logging	113-115	30-32, 230
	Mining	211-213	40-50
Construction	Construction	23	60
Durable Manufacturing	Wood Products	321	231-241
	Non-Metallic Minerals	327	250-262
	Primary and Fabricated Metals	331-332	270-301
	Machinery	333	310-332, 380
	Computers and Elec. Equipment	334-335	322-350, 371, 381
	Transportation Equipment	336	351-370
	Furniture and Related	337	242
	Misc. Durable Mfg.	339	372, 390-391
Non-durable Manufacturing	Food/Beverage/Tobacco	311-312	100-130
	Textiles, Apparel and Related	313-316	132-152, 220-222
	Paper Products	322	160-162
	Printing and Publishing	323, 511, 516	171-172
	Petroleum and Coal	324	200-201
	Chemical Products	325	180-192
	Plastics	326	210-212
Trade, Transportation and Utilities	Utilities	22, 486, 562	422, 450-472
	Wholesale Trade	42	500-571
	Retail Trade	44,45,722	580-691
	Rail, Air, Water Transportation	481-483	400, 420-421
	Other Transportation	484-492	401-410, 432
	Warehousing and Storage	493	411
Information	Information	515-517	440-442
Professional and Business Services	Professional and Business Services	518-519, 532-561	12, 20, 721-742, 801, 841, 852, 882-893
Financial Activities	Banks and Credit Intermediation	521-522	700-702
	Securities, Commodities, etc.	523-525	710
	Insurance and related	524	711
	Real Estate	531	712
Education/Health Services	Educational Services	61	842-851, 860
	Ambulatory Health Care	621	812, 820-830, 840
	Hospitals, Nursing, Residential Care	622-624	831-832, 861-871
Leisure/ Accommodation	Leisure/ Accommodation	512, 711-721	762-770, 800, 802, 810, 872
Other Services	Other Services	81	750-761, 771-791, 873-881

Table 2.11: Detailed Abstract Occupational Codes

Occupation Group	*occ1990dd* codes
Managers and Management Related	4-37
Engineers	43-65
Computer Scientists and Programmers	64,229,233
Scientists	65-83
Medical Occupations	84-106
Teachers	154-165
Lawyers/Social Science/Social Work	166-178
Sales	235-258
Technicians	203-228,234-235
Other (arts/media/entertainment/public safety)	183-199, 417-423

As my focus is on private non-farm sector, I extract the household sector from industry output from the BEA. This is done by subtracting the compensation of private households from both the value added and compensation of its corresponding sector - other services, excluding government. The remaining component of private household value added - the rental value of owner-occupied housing - is subtracted from value added in the real estate sector. Also, from the CPS data, there is an extra industry of unclassified manufacturing (1990 Census Code 392), which does not have a comparable analog in the BEA data. As such, I take the income and hours of this industry and allocate them proportionately across the other manufacturing industries.

The detailed occupation codes for abstract industries are developed from the consistent occupational codes of Autor and Dorn (2013). The following table, Table 2.11, lists each occupation group and the 1990 Census Occupational codes as used in the variable *occ1990dd*, documented in Dorn (2009).

Chapter 3

Acyclical Labor Productivity: It's Not as Simple as You Think

3.1 Introduction

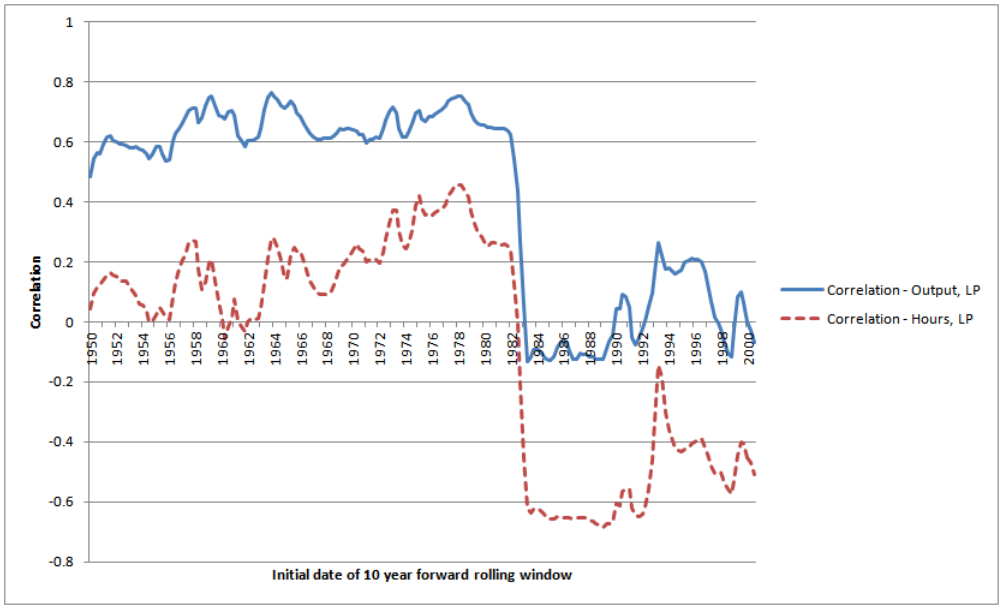
A recent literature documents that around the 1980s, labor productivity - measured as total output divided by total hours - turned from being strongly procyclical to being essentially acyclical.[1] Figure 3.1 displays this fact by plotting the 10 year rolling correlation between cyclical labor productivity and output, as well as between labor productivity and hours.

This fact is of interest for two reasons. First, the cyclical dynamics of productivity are central to theories of the business cycle, and changes in those dynamics could have strong implications for sorting between these theories. For example, traditional RBC models driven by Solow-residual technology shocks imply a strong positive correlation between output and labor productivity, but it is known that government spending shocks or leisure shocks can potentially undo this correlation, depending on the relative magnitude and volatility of these shocks.[2] The recent change in

[1] See Stiroh (2009), Gali and Gambetti (2009), Barnichon (2010), Gali and Van Rens (2010), and Garin et al. (2013).

[2] See Hansen and Wright (1992) for an early discussion of productivity and output response to a variety of shocks in the RBC model

Figure 3.1: Aggregate 10 Year Forward Looking Rolling Correlations Between Labor Productivity and Output/Hours



Data is from BLS Labor Productivity and Costs Series for the Nonfarm Business Sector. Labor productivity is defined as output per hour. All series have been logged and detrended using an HP filter with parameter 1600.

the cyclicality of labor productivity could provide valuable insight into the underlying drivers of the business cycle and how those have changed over time and how we should think about recent business cycle fluctuations.

The second reason that the recent acyclicality of labor productivity is of interest is because it is not the only feature of the business cycle to have changed around the 1980s. Perhaps the best known difference in the business cycle is the substantial decline in GDP volatility beginning in the 1980s, commonly known as the Great Moderation.[3] The fact that the cyclical dynamics of labor productivity have changed at about the same time as the decline of GDP volatility suggests that there might be a connection between these two events. Thus, the acyclicality of labor productivity may give insight into some of the underlying causes of the Great Moderation, a point made in Gali and Gambetti (2009) and Gali and van Rens (2010). Further, all three recessions since the 1980s have been characterized by so-called "jobless recoveries" and it is possible that the acyclicality of

[3]See Stock and Watson (2003) for a review of this fact and an early survey of the literature surrounding it.

labor productivity is related to this phenomenon as well, something suggested in Berger (2012) and Garin et al. (2013). The timing of the acyclicality of labor productivity links it to these recent business cycle changes and suggests that it might provide additional evidence towards understanding these changes.

The most common story to emerge in explaining the recent acyclical of labor productivity has been that some kind of labor frictions have changed or have become relatively more important, and a number of models have been able to replicate acyclical productivity hinging on abstract frictions.[4] However, it is unclear exactly which labor frictions in reality explain the change in the cyclicality in productivity. Suggested frictions include changes in unionization rates, changes in the prevalence of performance for pay practices, an increased importance of sectoral reallocation, and a decline in hiring costs due to the rise of online recruiting efforts. Thus, while it is possible to generate a change in the cyclicality of productivity by placing added emphasis on labor frictions in a relatively broad class of models, there is yet to emerge strong empirical evidence in favor of any one friction in particular.

In this chapter, I aim to deepen our understanding of acyclical productivity by providing additional empirical evidence and perspective on this recent change that can potentially illuminate its underlying causes. My analysis yields four key findings:

1. Acyclical productivity is not statistically synonymous with recent jobless recoveries.

2. Acyclical productivity is statistically synonymous with the change in the relative volatility of hours and output over the business cycle.

3. There is a great deal of heterogeneity in this relative volatility of hours and output across industries and across time.

4. Aggregate changes in the relative volatility of hours and output are also partially explained

[4]This is the basic story for Barnichon (2010), Gali and van Rens (2010), Nucci and Riggi (2011), Berger (2012), and Garin et al. (2013)

by changes in between industry co-movements.

Holding up these four points against the existing explanations for acyclical labor productivity, particularly the evidence from the industry level, the current stories based on changing labor frictions seem inadequate to explain a number of the industry-level central features generating aggregate acyclical productivity. With the current literature unable to reconcile these changes, acyclical productivity remains a puzzle and an even more complicated one at that.

Section 2 gives added documentation regarding the change in the cyclicality of labor productivity over the business cycle. Section 3 presents evidence showing the relationships between acyclical productivity and the aggregate phenomena of jobless recoveries and volatility changes. Section 4 considers industry level evidence on volatility changes and the role of industry co-movements in explaining acyclical productivity. Section 5 summarizes these conclusions and examines their implications for existing explanations of acyclical productivity, as well as considering possible future directions for this research. Section 6 concludes.

3.2 The Emergence of Acyclical Labor Productivity

In this section, I present the stylized facts regarding changes in the cyclical dynamics of aggregate labor productivity. For aggregate data, I use quarterly data on real output, hours and labor productivity (total real output divided by total hours) from the BLS Labor Productivity and Costs Series for the non-farm business sector covering 1947Q1 - 2013Q2. All variables are logged and detrended with an HP filter with parameter 1600. Additionally, consistent with the suggestion of Baxter and King (1999), I exclude the first and the last three years of the sample to avoid potential endpoint problems with the HP filter.

As stated in the introduction, since the 1980s, there has been a sharp reversal in the co-movement patterns of labor productivity with aggregate quantities, such as output and hours. Fig-

ure 3.1 plots the 10 year forward rolling correlation between real output and labor productivity, as well as the rolling correlation between hours and productivity. What immediately stands out is a sudden change in these correlations occurring between the rolling windows of 1982 Q2 - 1991Q1 and 1983 Q2 - 2010 Q2. Table 3.1 uses the latter of these window start points as a break point and reports the correlations from 1950 Q1 - 1983 Q1 and from 1983 Q2 until 2010 Q2[5]. In the early period, the correlations of output and hours with productivity were 0.65 and 0.21, respectively. However, in the subsequent period, these correlations have dropped to -0.04 and -0.52, respectively, a net change of between -0.65 and -0.7 for each series.[6] As evidenced in both Figure 3.1 and Table 3.1, the drop in correlations between the rolling windows of 1982 Q1 - 1991 Q4 and 1983 Q2 - 1993 Q1 is even larger, falling from 0.62 and 0.25 to -0.13 and -0.62 for output and hours respectively. And although these correlations rose slightly in the 1990s and early 2000s, they did not return to their pre-1980s levels, and this rise appears to have been only temporary.

This striking fact has been well documented in Barnichon (2010), Gali and van Rens (2010) and Garin et al. (2013). They also document that these same qualitative relationships hold for using employment and output per worker instead of hours and output per hour, and document that there have been changes in correlations with productivity and other quantities, such as vacancies and unemployment. In what follows, I restrict my attention to the correlation between productivity and output, however the results I obtain will apply to the correlation between productivity and other aggregate quantities as well.

A number of recent papers have attempted to explain this sudden recent change in the cyclical behavior of labor productivity, with a great deal of emphasis placed on the role of labor frictions. Barnichon (2010) proposes a model in which a changing composition of non-technology and technology shocks (suggested to be owing to improved monetary policy in the Great Moderation), coupled with a fall in hiring costs can generate a change in the cyclical nature of productivity such that productivity and unemployment become positively correlated after the 1980s. The idea

[5]However, all results are robust to varying the break date throughout the early 1980s

[6]Hagedorn and Manovskii (2011) have shown that when you instead use hours from household surveys, the following results are qualitatively the same, though quantitatively different.

Table 3.1: Correlations Between Productivity and Output or Hours Across Time Periods

	Output, Productivity	Hours, Productivity
1950 Q1 - 1983 Q1	0.6474	0.2137
1983 Q2 - 2010 Q2	-0.0373	-0.5175
Difference	-0.6847	-0.7312
1982 Q1 - 1991 Q4	0.6229	0.2415
1983 Q2 - 1993 Q1	-0.1307	-0.6235
Difference	-0.7536	-0.8650

Data is from BLS Labor Productivity and Costs Series for the Nonfarm Business Sector. Labor productivity is defined as output per hour. All series have been logged and detrended using an HP filter with parameter 1600.

is that an increased shift toward shocks generating a positive correlation between unemployment and productivity (in this case, productivity shocks because of nominal rigidities), coupled with reduced utilization of the effort margin in favor of the hours/employment margin (because of cheaper hiring costs and/or an increased elasticity of hours per worker), leads to labor productivity being high in times of high unemployment. Gali and van Rens (2010) also allow for technology and non-technology shocks, but focus strictly on changes in labor frictions which change utilization of the effort margin relative to the hours margin. Nucci and Riggi (2011) follow in this line of research, but suggest that increased adjustment in hours instead of effort could stem from an increased cost in effort compensation instead of a decreased cost in hours adjustment. In different frameworks, Berger (2012) and Garin et al. (2013) relate the recent acyclicality of productivity to jobless recoveries and suggest that these two could be linked if either firms have gained the ability to selectively fire unproductive workers, or if a change in the composition of shocks has favored shocks inducing frictional labor reallocation, respectively.

At the heart of each of these stories is a key labor friction that has either changed or become more prominent through some change in the composition of shocks. Barnichon (2010), Gali and van Rens (2010), and Berger (2012) suggest some increased flexibility of the labor market, proposing changes such as the decline of unionization or the rise in online recruiting resources as explaining the more flexible labor market. Flexibility could also come in compensation, as suggested by Nucci and Riggi (2011), and they suggest that the increased prevalence of performance for pay

practices could be the driver here. Garin et al. (2013) consider an increased importance of reallocation frictions as the driving force, and this could be consistent with ongoing structural changes with the decline of manufacturing, or even reallocation across different occupations observed in the past several decades. Thus, within explanations relying on labor frictions, there are yet many plausible mechanisms which could be responsible for these changes.

3.3 Acyclical Productivity and Aggregate Changes

The cyclicality of productivity was not the only thing to change in the 1980s. In all the recessions since the 1982 recession, recoveries have been marked by unusually sluggish growth in employment and hours, especially relative to output, giving rise to so-called "jobless recoveries." And beginning with the end of the 1982 recession, we observe a marked decline in the volatility of GDP and a wide variety of other aggregate quantities, beginning a period known as the Great Moderation. In this section, I show that the former of these two phenomena - jobless recoveries - is not synonymous with acyclical productivity. That is, we observe a decline in the correlation between output and productivity across all phases of the business cycle, not just in recoveries. Thus, while there may yet be some relationship between the underlying causes of both jobless recoveries and acyclical productivity, the two do not completely coincide. On the other hand, I show that changes in volatilities in output and hours occurring with the Great Moderation directly map into the correlation between productivity and output and these changes mechanically explain the change in the cyclicality of labor productivity.

3.3.1 POINT #1: Acyclical productivity is not statistically synonymous with jobless recoveries

The notion of a jobless recovery, as witnessed in the past several recessions, is that growth in hours and employment have been quite sluggish relative to past recoveries, whereas the recovery in output has been slow, but closer to historic recoveries than hours. This definition, in and of itself, mechanically generates countercyclical movements in labor productivity. During a recovery,

Table 3.2: Correlation Between Output and Productivity Beginning in 1982 Q2 and 1983 Q3, with Varying Forward Rolling Window Lengths

	1982 Q2	1983 Q3	**Difference**
10 years	0.62	-0.13	**-0.75**
8 years	0.66	0.00	**-0.66**
6 years	0.81	-0.22	**-1.04**
4 years	0.82	-0.11	**-0.93**

Data is from BLS Labor Productivity and Costs Series for the Nonfarm Business Sector. Labor productivity is defined as output per hour. All series have been logged and detrended using an HP filter with parameter 1600.

output is still below trend, but in jobless recoveries, it is growing faster than hours, leading to a rise in output per hour. Further, using a ten year rolling window for the correlation between output and productivity, a break in this relationship around 1982 could simply be because the ten year window intersects the jobless recovery first observed with the 1990-1991 recession. Therefore, one possible conjecture regarding acyclical labor productivity is that it is statistically synonymous with jobless recoveries, a mere reflection of these changes in the behavior of recoveries. Tables 3.2 and 3.3, however, provide evidence that there is more to acyclical productivity than simply jobless recoveries.

Table 3.2 demonstrates that the sharp decline in the rolling correlation between output and productivity occurring around 1982 is not sensitive to the length of rolling window, and thus even excluding observations from the 1990s recession, there is still large drop in this correlation around 1982. Table 3.2 gives the value for the rolling correlation starting in 1982 Q2 and 1983 Q3 under window lengths of 4, 6, 8 and 10 years. For the window lengths of 4 and 6 years, no observations intersect the 1990-1991 recession or its recovery. In particular, looking at a six year rolling window, the correlations in the rolling windows beginning in 1982 Q2 and 1983 Q3 is 0.81 and -0.22, reflecting a drop of 1.04 points in this short time horizon. Using a four year rolling window, the drop is of a similar magnitude, with the correlation falling by 0.93. These declines are actually even larger than those observed using 8 or 10 year windows which intersect the 1990s recession and recovery. Acyclical productivity was already occurring before the first jobless recovery.

Table 3.3 examines how the correlation between output and productivity has changed across

Table 3.3: Output-Productivity Correlations Across Different Business Cycle Phases, Before and After 1982 Q2

	1950 Q1 - 1983 Q1	1983 Q2 - 2010 Q2	Difference
All Phases	0.6474	-0.0373	-0.6847
Recession	0.6623	0.3142	-0.3481
Recovery	0.7351	0.4490	-0.2861
Expansion	0.5992	0.1451	-0.4541

Data is from BLS Labor Productivity and Costs Series for the Nonfarm Business Sector. Labor productivity is defined as output per hour. All series have been logged and detrended using an HP filter with parameter 1600. Recession time periods are defined by the NBER Business Cycle dates. Recovery time periods are defined as the continuous period of time following a recession that hours remain below trend. Expansion time periods are the residual of recession and recovery time periods.

different phases of the business cycle. I consider three phases of the business cycle - expansion, recession and recovery. I use the NBER recession dates to define the recession phase, and define a recovery as the continuous period following a recession during which hours remains below trend. Expansions are then defined as the residual of these two phases. I focus on the recovery behavior of hours to highlight jobless recoveries in particular, as the recent recoveries have witnessed longer recovery times for hours than output.

Table 3.3 reports the correlation between output and productivity conditional on being in each of these phases of the business cycle in the period prior to 1983 and the period following 1983. The correlation between output and productivity has declined in all phases of the business cycle. In each of the phases of recession, recovery and expansion, the correlation between output and productivity was quite high in the early period, between 0.6 and 0.7. However, since then, the correlation has declined in each of these three phases, although it has declined the least in the recovery stage of the business cycle, falling from 0.74 to 0.45, and declined the most in the expansion stage, falling from 0.6 to 0.15.

Using different definitions of what constitutes a recovery phase, such as the period after a recession before hours returns to its pre-recession level, or using employment instead of hours, can lead to a slightly larger measured decline in correlation during the recovery phase. However the decline in correlation during the expansion phase is always at least as great as the decline in the

recovery phase, regardless of the recovery definition used.

Figure 3.2 further reinforces this point by plotting output and productivity throughout the post-war period. In the period prior to the 1980s, labor productivity rose and fell fairly closely with output. However, after the 1982 recession, the two series begin to diverge. Beginning in the mid to late 1980s, output begins to rise in the boom preceding the 1990s recession, but productivity drops below zero, hovering around there until just after the initial fall in output, when it jumps up while output remains well below trend. Both output and productivity hover around zero until the run up to the 2001 recession, when output expands considerably while productivity continues to hover around zero. Again, just after output crashes, productivity spikes up. And then productivity has been falling since 2003 as output recovered, and only surged back up again in the trough of the Great Recession.

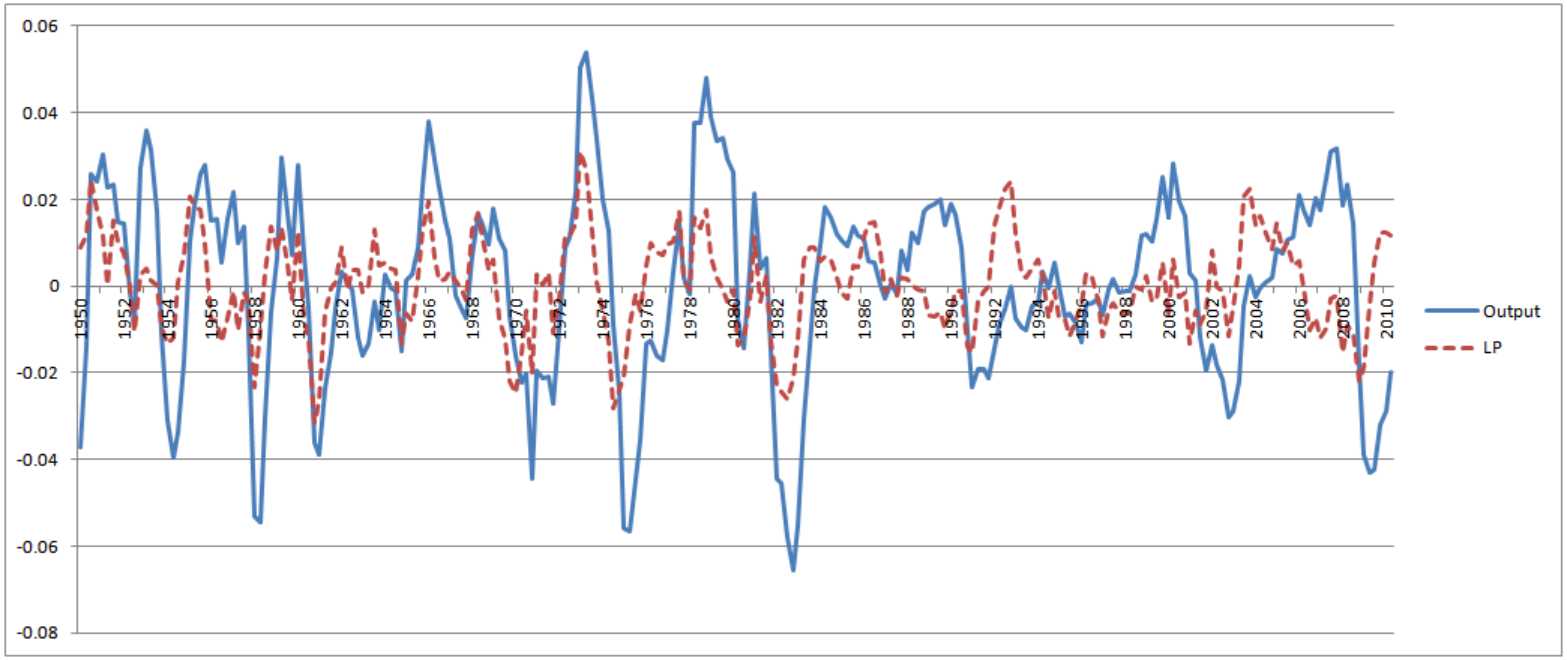
Acyclical productivity is occurring across all phases of the business cycle, and not just during periods that might be termed jobless recoveries. Thus, although there may yet be some joint causes for acyclical productivity and jobless recoveries, but the two phenomenon are not one and the same.

3.3.2 POINT #2: Acyclical productivity is statistically synonymous with changes in relative volatilities occurring in the Great Moderation

In the 1980s, at the same time as the cyclical dynamics of productivity changed, the volatility of GDP and other aggregate quantities fell sharply and have remained low for a long period of time. This sharp and subsequently persistent decline of volatility has caused this period to be called the Great Moderation. Given the exact coincident timing of these two events, papers such as Stiroh (2009) and Gali and Gambetti (2009) have tried to understand some of the causes of the Great Moderation using the additional evidence of acyclical productivity.

How exactly might acyclical productivity reflect changes in volatilities occurring with the Great Moderation? Fundamentally, the dynamics of labor productivity simply describe how output and hours relate over the business cycle - the log of labor productivity is simply the difference between the log of output and hours. Thus, it is instinctive to think of changes in the dynamics of labor

Figure 3.2: Log Deviations from Trend: Output and Labor Productivity



Data is from BLS Labor Productivity and Costs Series for the Nonfarm Business Sector. Labor productivity is defined as output per hour. All series have been logged and detrended using an HP filter with parameter 1600.

productivity to be changes in either the co-movement of hours and output or changes in their relative volatilities. The following lemma formalizes this notion.

Lemma: The correlation of output with labor productivity, $\rho_{y,y-h}$, can be written purely as a function of the relative volatility between output and hours, $\frac{\sigma_h}{\sigma_y}$, and the correlation between output and hours, $\rho_{y,h}$. Further, $\rho_{y,y-h}$ is weakly decreasing in $\frac{\sigma_h}{\sigma_y}$.

Proof: See the Appendix.

The result of this lemma is that we can write the correlation of output and productivity in the following formula:

$$\rho_{y,y-h} = f\left(\frac{\sigma_h}{\sigma_y}, \rho_{y,h}\right) \equiv \frac{1 - \frac{\sigma_h}{\sigma_y}\rho_{y,h}}{\sqrt{1 + \left(\frac{\sigma_h}{\sigma_y}\right)^2 - 2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}}}$$

The intuition for this result is straightforward. Consider the limiting case in which output and hours are perfectly correlated, $\rho_{y,h} = 1$. If output is more volatile than hours, then its correlation with labor productivity is exactly 1; if the opposite - output is less volatile than hours - then their correlation is exactly -1. Thus, in the first case where output is more volatile than hours, in times of expansion, output, hours and productivity all rise, as output rises more than hours. However, if hours are more volatile than output, then in expansionary times, while output and hours both rise, productivity falls because hours have risen more than output. Similarly, if output and hours were perfectly negatively correlated, $\rho_{y,h} = -1$, then these results would be the exact opposite, with the correlation between output and productivity being exactly 1 if hours are more volatile than output and exactly -1 if hours are less volatile than output.

So what has changed in terms of volatilities and co-movements between output and hours in the past thirty years? The left panel of Figure 3.3 plots the relative rolling standard deviations of hours and output for the postwar period, as well as the rolling correlation between output and hours over the same time horizon and Table 3.4 presents these numbers for the the pre-1983 and post-1983 periods. Hours and output are consistently highly correlated throughout the entire sample, and this correlation shows virtually no change between the pre and post 1983 periods, as shown

Table 3.4: Changes in Components of Correlation of Labor Productivity with Output

	σ_h/σ_y	$\rho_{y,h}$
1950 Q1 - 1983 Q1	0.780	0.883
1983 Q2 - 2010 Q3	1.168	0.874
%Δ	+49.7%	-1.0%
1982 Q1 - 1991 Q4	0.806	0.910
1983 Q2 - 1993 Q1	1.268	0.857
%Δ	+57.3%	-5.8%

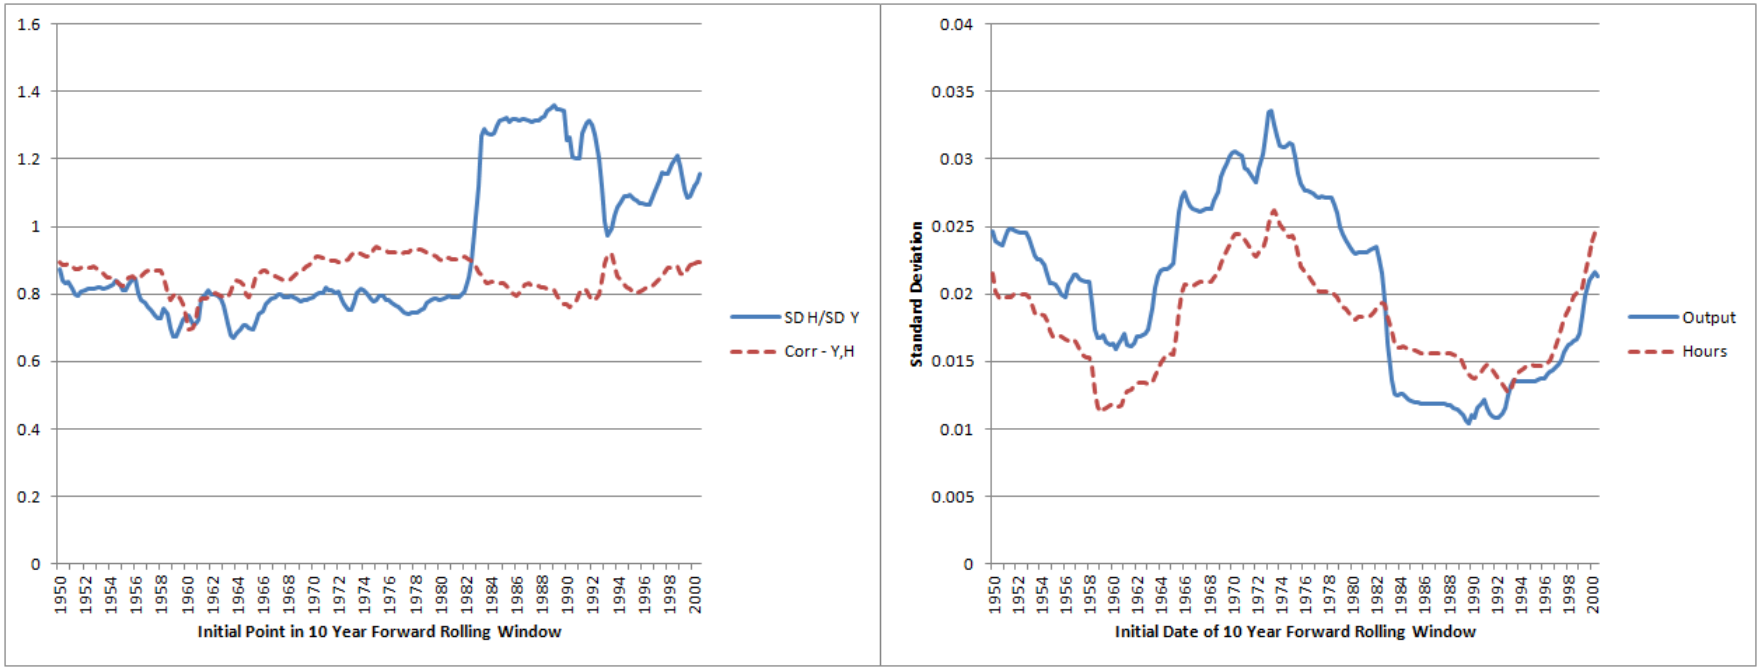
Data is from BLS Labor Productivity and Costs Series for the Nonfarm Business Sector. Labor productivity is defined as output per hour. All series have been logged and detrended using an HP filter with parameter 1600.

in Table 3.4. On the other hand, the relative volatility of hours and output changed dramatically in the 1980s, and in Figure 3.3, looks almost like the mirror image of the correlation of labor productivity and output over this time period. The standard deviation of hours was about 80% that of the standard deviation of output in the period prior to the 1980s. However, around the 1980s, the standard deviation of hours jumps up suddenly to almost 130% the volatility of output and although it has come down since the 1990s, it still is about 110-120% that of the standard deviation of output. For the post-1983 period, the volatility of hours is 117% that of the volatility of output, a 50% increase from the relative volatility in the period prior to 1983.

The right hand panel of Figure 3.3 shows the separate rolling standard deviation series for hours and output. As is well known, the volatility of output declines substantially in the early 1980s and achieves a postwar low level of volatility until near the very end of the sample, when the Great Recession begins. The volatility of hours closely mirrors the volatility of output, but does not decline nearly as much as the volatility of output around the 1980s, and has since been more volatile than output. This change in the relative volatility of hours and output in the Great Moderation was first observed by Gali and Gambetti (2009) and has been given subsequent attention in Gali and van Rens (2010).

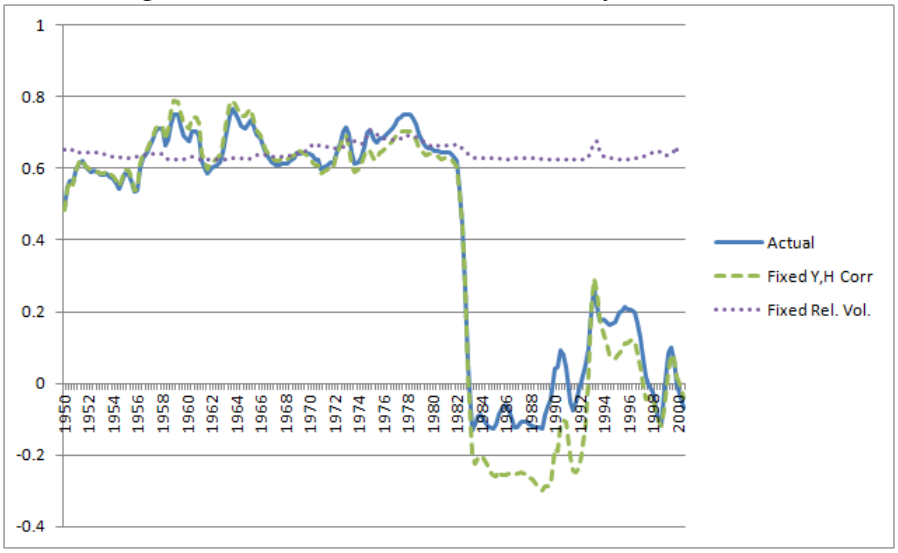
This evidence makes it clear that the sudden change in the relative volatility of hours and output occurring in the Great Moderation is precisely the reason why productivity correlations have changed so much. To further support this point, Figure 3.4 plots two counterfactual rolling corre-

Figure 3.3: Aggregate 10 Year Forward Looking Rolling Standard Deviations of Output and Hours



Data is from BLS Labor Productivity and Costs Series for the Nonfarm Business Sector. All series have been logged and detrended using an HP filter with parameter 1600.

Figure 3.4: Counterfactual Productivity Correlations



Data is from BLS Labor Productivity and Costs Series for the Nonfarm Business Sector. Labor productivity is defined as output per hour. All series have been logged and detrended using an HP filter with parameter 1600. The green-dashed series represents fluctuations in the correlation between output and productivity where the correlation between output and hours is fixed throughout the entire sample at its level for the period 1950 Q1 - 1983 Q2. The purple dotted series represents fluctuations in the correlation between output and productivity where the relative volatility of hours and output is fixed throughout the entire sample at its level for the period 1950 Q1 - 1983 Q2.

lations of output and productivity. A counterfactual correlation is given by fixing one input into the above function f at its pre-1983 level, while allowing the other input to vary throughout the sample. While these "counterfactual" plots do not necessarily have a complete economic intuition, they highlight the quantitative relevance of the two channels affecting the cyclicality of labor productivity - the relative volatility of hours and output and the correlation between the two.

In Figure 3.4, the plot representing the actual correlation is virtually indistinguishable from the plot where the output and hours correlation is held fixed. And any discrepancies between the two only serve to suggest that if anything, the subtle movements in the output and hours correlation are slightly increasing the correlation between output and productivity. And in the counterfactual correlation series where relative volatility is held fixed, there are virtually no fluctuations throughout the entire postwar period.

All this evidence has been presented to make this simple point - the recent acyclicality of

productivity and the change in the relative volatility of hours and output are the exact same stylized fact. This stands in contrast to Gali and van Rens (2010), who present these two facts are two separate pieces of evidence in favor of the theory they are considering. As the changes in relative volatility of hours and output are synonymous with the recent acyclicality of productivity, I focus my attention on the relative volatility statistic for the remainder of the chapter.

3.4 Industry Level Changes

To this point, I have focused on changes in the behavior of aggregate labor productivity and demonstrated its relation to other aggregate changes. In particular, the recent changes in the dynamics of aggregate labor productivity are synonymous with the recent changes in the relative volatilities of hours and output. In considering possible explanations for this change, it is helpful to know if these changes have been occurring evenly across different industries or if they have been disproportionately concentrated in a particular subset of industries. It is also useful to understand how changes within industries have evolved over time. The industrial composition of economic activity has changed substantially over time and although the changes in the behavior of productivity appear to occur quite suddenly, it is possible that there are additionally changes in the underlying industry composition that are partially responsible for this change as well.

For example, one possible explanation for these changes is that declines in unionization have loosened labor frictions, leading to more adjustment in hours and employment relative to effort. If this is the case, then we might expect to see changes in the behavior of productivity to be particularly prominent in industries where there has been the most change in unionization rates. On the other hand, if the reason for the change in the cyclicality of productivity is some change in the value of leisure common to all individuals, we should expect to see very uniform behavior of productivity across industries. Further, given that the change in productivity dynamics seems concentrated around the 1980s, many possible explanations are tested by examining specific changes occurring in the 1980s and 1990s, but if there are changes across time at the industry level, this may potentially open up the set of feasible explanations beyond those with stark changes occurring

in the 1980s.

In considering heterogeneity at the industry level, I take advantage of the fact that the recent acyclicality of productivity is synonymous with the changes in the relative volatility of hours and output and focus on industry level changes in the relative volatility of output and hours. Focusing on relative volatility instead of productivity correlations is beneficial because it is often easier to think intuitively about volatilities than correlations, volatility statistics have nice linear properties making total decomposition easier, and focusing on industry level volatility allows for a greater connection to the Great Moderation literature which has extensively studied volatility changes at the industry level.[7]

For industry level data, I use annual data on hours and output from the Groningen Growth and Development Centre's Productivity Database for eight distinct non-overlapping sectors of the economy from 1950-2007.[8] The covered sectors are mining, manufacturing, utilities, wholesale and retail trade, construction, financial/ real estate / professional/business services (FIRE+), transportation / communication / storage services (TCS), and personal/social/community services (PSC).[9] While other industry level data sets provide potentially more industries, I know of no other data set that covers a comparably long time window, has data spanning the total economy, and has actual data on hours. As with the aggregate data, all series are logged and filtered with the HP filter, here using the parameter of 6.25 for annual data, as suggested by Ravn and Uhlig (2002) and I have again omitted the first and last three years of the sample from my analyses, consistent with the suggestions of Baxter and King (1999).

[7]See, for example, Warnock and Warnock (2000) and Stiroh (2009).

[8]The original data set only extends through 2005, but I use data from the EUKLEMS database, which gathers from the same sources under the same classification scheme, to extend the data the last two years. Further, the aggregate patterns of volatility and correlation are exactly observed in the GGDC data; I have omitted these results for brevity's sake, but they are available by request.

[9]These sector definitions are according to International Standard Industrial Classification (ISIC) codes, revision 3.

Table 3.5: Industry Level Relative Volatility Before and After 1983

	1953-1982	1983-2004	% Δ
Mining	1.22	1.22	**0%**
Manufacturing	0.77	0.84	**+9%**
Construction	0.85	0.86	**+2%**
Utilities	0.48	0.33	**-31%**
Trade	0.43	0.57	**+31%**
FIRE+	1.03	1.72	**+67%**
TSC	0.66	1.15	**+74%**
PCS	0.63	0.44	**-30%**

Data comes from the GGDC 10 Sector Database; industry classifications follow ISIC, rev. 3. Underlying data has been logged and filtered with an HP filter with parameter 6.25, consistent with Ravn and Uhlig (2002).

3.4.1 POINT #3: There is a great deal of heterogeneity in the relative volatility of hours and output across industries and across time.

Table 3.5 presents the relative volatility of hours and output in each of the eight industries for the 1953-1982 and 1983-2004 periods. The first clear observation from Table 3.5 is that there is a great deal of heterogeneity in this relative volatility across industries, even before 1983. In the period from 1953-1982, the relative volatility of hours and output was only 0.43 in the wholesale and retail trade industries, but as large as 1.03 in the FIRE+ industries and 1.22 in the mining industry. However, aside from the mining and FIRE+ industries, the relative volatilities in all industries were less than 1 - that is, output was more volatile than hours in almost all industries prior to 1983.

Even more importantly, there is also substantial heterogeneity in the changes in relative volatility across industries between these two periods. Three industries saw substantial increases in relative volatility (wholesale and retail trade, FIRE+, and transportation/ storage/ communication services), two industries saw substantial declines in relative volatility (utilities and personal/ community/ social services), and three industries saw minimal changes in relative volatility (mining, manufacturing and construction). Generally speaking, changes in relative volatility, either increases or decreases, have come in the services industries. Of the three "goods" industries -

mining, manufacturing and construction - there is only a 9% increase in relative volatility in manufacturing, a 2% increase in construction, and no change in the mining industry. On the other hand, each of the services industries saw at least a 30% change in relative volatility, either increasing or decreasing, between these two sample periods. And it is in the larger services industries - trade, FIRE+ and transportation/storage/communication services - where there have been increases in relative volatility, with increases of 31%, 67%, and 74% in all three industries, respectively.

This heterogeneity in changes in relative volatility has potentially significant implications for the validity of explanations regarding acyclical productivity. Although there has been an aggregate increase in the volatility of hours relative to output, this increase is only observed in three of these eight major subsectors and there have very been minimal changes in relative volatility for all goods producing sectors of the economy. For example, an explanation of acyclical productivity relying on sharp declines in unionization seems to not fit the facts here. As described in Hirsch (2008), unionization and unionization declines have been greatest in industries such as mining, manufacturing, construction, utilities, and transportation services. However, there have been minimal changes in relative volatility in mining, manufacturing and construction, and even a decrease in relative volatility in the utilities industry. There has still been a substantial increase in relative volatility in transportation/storage/communication industries, but the fact that other industries with substantial unionization declines have not seen much change in relative volatility casts a substantial amount of doubt on this particular hypothesis. I return to this point and consider other implications of this evidence in Section 5.

Given that the aggregate change in relative volatility appears to have been very sudden around the 1980s and not gradual over time, it is also insightful to examine whether the observed changes in relative volatility at the industry level have occurred in a similar fashion, or whether they have been more gradual over time. Figure 3.5 presents relative volatility in a ten year rolling window for each of these industries across the sample time frame, with a line at the year 1983, when the sudden change in the aggregate series is observed. The first panel of Figure 3.5 shows relative volatility for the goods industries, where the split sample comparison showed minimal change in the relative

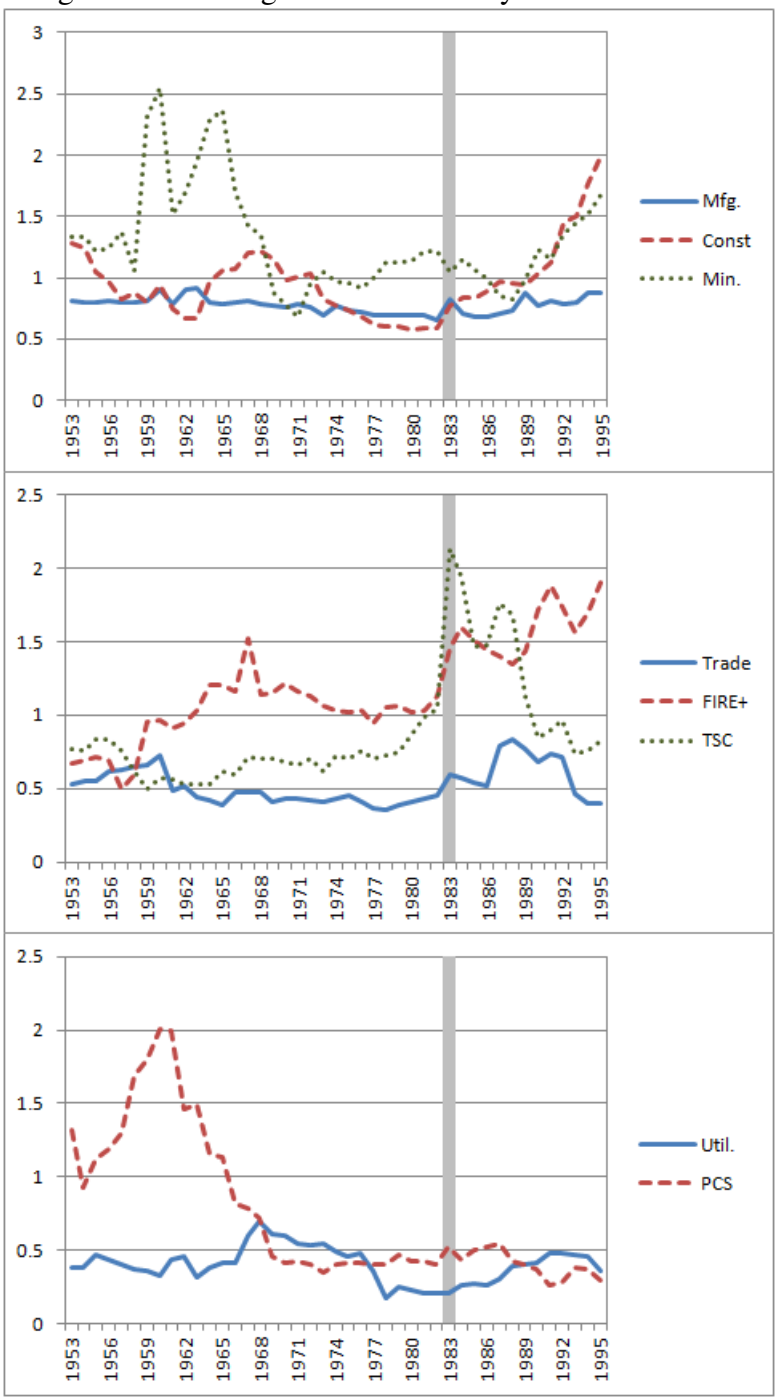
volatility of hours to output. What we observe is that although there appear to have been minimal changes between the broader periods of before and after 1983, there has actually been quite a bit of change over time. Hours were quite volatile relative to output in both the construction and mining industries in the 1960s, but relative volatility declined in the 1970s and remained fairly low for a long period of time. More recently, relative volatility has been rising again for these two industries since the 1990s, particularly in construction, where the volatility of hours is nearly double that of output at the end of the sample window, a postwar high in relative volatility. On the other hand, relative volatility in manufacturing has been fairly steady throughout time, though it does appear to be gradually rising since the 1980s.

Relative volatility in the trade, FIRE+ and TSC industries, as shown in the second panel of Figure 3.5, mirrors the aggregate behavior much more closely. From the early 1960s through the 1970s, relative volatility was fairly constant for each of these three industries. Then, in the 1980s, relative volatility increased substantially, particularly around 1983, but also before, in the case of the TSC industries, and after, in the case of the trade industries. However, there is where the similarity between these industries' and the aggregate relative volatility begin to differ. For the trade and TSC industries, after the 1980s, relative volatility declined, and quite substantially. In both trade and TSC, the level of relative volatility returned to levels comparable to the pre-1980s period by the end of the sample. Relative volatility in FIRE+ industries, however, has only continued to grow since the 1980s and shows no sign of decline.

The last panel of Figure 3.5 shows the relative volatility in the utilities and PCS industries. While there were periods of very high relative volatility in the 1950s and 1960s in these industries, since then, there has not very much movement in their relative volatilities. Relative volatility in the utilities industry increased gradually throughout the 1980s and 1990s, but such changes have only been a return to slightly higher levels in the pre-1980s period. Neither series shows any meaningful "jump" around the 1980s.

In summary, although the aggregate changes in these series appear to have been primarily concentrated in the early 1980s, there has been substantial ongoing change ever since the 1980s at the

Figure 3.5: Rolling Relative Volatility Across Industries



Data comes from the GGDC 10 Sector Database; industry classifications follow ISIC, rev. 3. Underlying data has been logged and filtered with an HP filter with parameter 6.25, consistent with Ravn and Uhlig (2002).

industry level. The trade, FIRE+ and TSC service industry mimic the aggregate change with a substantial jump in relative volatility around the 1980s, but have not been constant since, with relative volatility continuing to increase in FIRE+ industries, and relative volatility decreasing to pre-1980s levels in trade and TSC. Meanwhile, relative volatility has recently started to increase substantially in the 1990s and 2000s for the mining and construction industries, with gradual increases also occurring in manufacturing and utilities industries.

How does this contribute to explaining acyclical productivity and aggregate relative volatility? It suggests that one-time permanent changes that occurred in the 1980s are not apt to explain well these aggregate changes. Although there were jumps in relative volatility around the early 1980s in some industries, not all of those changes have proved to have permanent effects, and some have only been the start of more gradual increases in relative volatility through the past three decades. Understanding this timing heterogeneity broadens the set of possible explanations for recent changes by removing the 1980s timing restriction. However, given the substantial heterogeneity occurring across industries and their timing, this also suggests that there may be a number of causes generating aggregate acyclical productivity and not just one sole factor which has changed. Additional implications of this evidence are discussed in Section 5.

3.4.2 POINT #4: Aggregate changes in the relative volatility of hours and output are also partially explained by changes in between industry co-movements.

Although there is much can be learned by examining variation in industry level relative volatilities, it is not immediately obvious how these industry level changes in relative volatility aggregate up to whole economy. In this section, I show how this aggregation occurs and show that there may be other industry level factors affecting the aggregate relative volatility changes outside of within industry changes in relative volatility. In particular, I highlight how changes in the co-movements between cross-industry hours and changes in the co-movements of outputs across industries have also played an important role in generating the change in aggregate relative volatility, and thus in

the cyclicality of labor productivity.

To understand how industry level relative volatilities aggregate, I first decompose aggregate volatility into industry volatilities. Consider the following first order approximation about some point $ln(\bar{Y})$ for log aggregate output (and similarly for hours as well):

$$ln(Y_t) \approx ln(\bar{Y}) + \sum \left(\frac{\bar{y}_i}{\bar{Y}}\right) ln(y_{it}) - \sum \left(\frac{\bar{y}_i}{\bar{Y}}\right) ln(\bar{y}_i)$$

Given that the HP filter is a linear operator, this approximation still holds if you replace all components with their deviations from trend. Further, given that the constant terms will not be relevant for volatility, this leaves the following expression for the decomposition of the variance of log aggregate output (where hats denote deviations from trend):

$$Var\left(\widehat{ln(Y_t)}\right) \approx Var\left(\sum \left(\frac{\bar{y}_i}{\bar{Y}}\right) \widehat{ln(y_{it})}\right) =$$
$$\sum \left(\frac{\bar{y}_i}{\bar{Y}}\right)^2 Var\left(\widehat{ln(y_{it})}\right) + 2\sum_{i<j} \left(\frac{\bar{y}_i}{\bar{Y}}\right)\left(\frac{\bar{y}_j}{\bar{Y}}\right) Cov\left(\widehat{ln(y_{it})}, \widehat{ln(y_{jt})}\right)$$

The relationship between aggregate output volatility and industry output volatilities is determined by three factors - the individual industry output volatilities, the industry shares of total output, and the covariance of outputs between industries. Thus, a change in aggregate output volatility can come about in three different ways. One, individual industry volatilities can change. Two, the industry composition of aggregate output can change in such a way to put more or less weight on particularly volatile industries. Three, the co-movement properties of outputs across industries can change, generating higher volatility through greater synchronization across the business cycle, or lower aggregate volatility through decreased synchronization or even negative correlations across industry output fluctuations.

The relationship between aggregate and industry relative volatility is very similar. Using the same decomposition for hours, aggregate relative volatility can be written in terms of industry level

volatilities, shares and covariances as follows:

$$\frac{Var\left(\widehat{ln(H)}\right)}{Var\left(\widehat{ln(Y)}\right)} = \frac{\sum_{i=1}^{n}\left(\frac{\bar{h}_i}{\bar{H}}\right)^2 Var\left(\widehat{ln(h_{it})}\right) + 2\sum_{i<j}\left(\frac{\bar{h}_i}{\bar{H}}\right)\left(\frac{\bar{h}_j}{\bar{H}}\right) Cov\left(\widehat{ln(h_{it})}, \widehat{ln(h_{jt})}\right)}{\sum_{i=1}^{n}\left(\frac{\bar{y}_i}{\bar{Y}}\right)^2 Var\left(\widehat{ln(y_{it})}\right) + 2\sum_{i<j}\left(\frac{\bar{y}_i}{\bar{Y}}\right)\left(\frac{\bar{y}_j}{\bar{Y}}\right) Cov\left(\widehat{ln(y_{it})}, \widehat{ln(y_{jt})}\right)}$$

This decomposition highlights a simple point - changes in aggregate relative volatility might also be explained by changes in industry level relative covariances. That is, the relative volatility of hours and output could increase if the covariances between industry hours increase relative to the covariances between industry outputs. Because changes in covariance implicitly embody changes in variance, I concentrate my attention on changes in relative correlations. So how have correlations between industry hours and correlations between industry outputs changed? Given that sector size matters for the aggregate contribution of changes in these correlations and for the sake of brevity, I restrict my attention to correlations between the five largest industries in the data - manufacturing, construction, trade, FIRE+, and TSC. Table 3.6 presents the correlation between hours across industries and the correlation between outputs across industries over the periods 1953-1982 and 1983-2004.

From Table 3.6, it is fairly clear to see that correlations in output fluctuations across industries have declined in the latter period, whereas correlations in hours fluctuations across industries have remained fairly constant or increased. On average, correlation between outputs across industries declined by -0.33 and correlation between hours across industries increased by 0.07. For each pair of industries, with the exception of FIRE+ and trade, the correlation between industry hours increased relative to the correlation between industry outputs. Thus, changes in these relative correlations appear to be occurring across most industry pairs and thus may have a significant role in explaining the change in aggregate relative volatility.

Again, what does this imply for acyclical productivity and aggregate relative volatility? This suggests that deeper changes in the entire network structure of the economy that may have occurred in the Great Moderation could be important for explaining the recent acyclicality of productivity.

Table 3.6: Between Industry Correlations in Hours and in Outputs Before and After 1983

	Output Correlations			Hours Correlations		
	1953-1982	1983-2004	**Diff**	1953-1982	1983-2004	**Diff**
Mfg, Const.	0.78	0.57	**-0.21**	0.80	0.85	**+0.05**
Mfg, Trade	0.83	0.20	**-0.63**	0.79	0.87	**+0.08**
Mfg, TSC	0.85	0.43	**-0.42**	0.86	0.90	**+0.04**
Mfg, FIRE+	0.52	0.54	**+0.02**	0.53	0.77	**+0.24**
Const., Trade	0.82	0.60	**-0.22**	0.88	0.85	**-0.03**
Const., TSC	0.75	0.04	**-0.71**	0.81	0.86	**+0.05**
Const., FIRE+	0.61	0.54	**-0.07**	0.79	0.83	**+0.04**
Trade, TSC	0.68	-0.04	**-0.72**	0.78	0.88	**+0.10**
Trade, FIRE+	0.40	0.48	**+0.08**	0.61	0.65	**+0.04**
TSC, FIRE+	0.76	0.37	**-0.39**	0.63	0.77	**+0.14**
AVERAGE	0.70	0.37	**-0.33**	0.75	0.82	**+0.07**

Data comes from the GGDC 10 Sector Database; industry classifications follow ISIC, rev. 3. Underlying data has been logged and filtered with an HP filter with parameter 6.25, consistent with Ravn and Uhlig (2002).

Given that the changes in co-movements between hours across industries have been minimal, while most of the change has occurred in the co-movements of outputs across industries, it seems particularly unlikely that stories centering around labor frictions are apt to explain these correlational changes.

At this point, I have merely presented evidence on changes on individual relative volatilities and individual relative correlations across industries, but not provided any evidence to suggest which of these changes are most important for explaining aggregate relative volatility. Naturally, the contributions of these changes will depend quite a bit on sector sizes and the actual levels of volatility and correlation. The Appendix contains a complete decomposition of aggregate relative volatility at the two sector level - goods and services - and quantifies the contribution of each component of the decomposition to changes in aggregate relative volatility between 1953-1982 and 1983-2004. The conclusion reached from such a decomposition is that 59% of the increase in aggregate relative volatility comes from changes in relative volatility in the services sector, 31% of the increase comes from changes in relative correlation, and the remaining change is explained by a combination relative volatility in manufacturing and composition factors. In short, the increase in

relative volatility in the services industries and the increase in hours correlations relative to output correlations seem central to understanding the recent acyclical behavior of labor productivity.

3.5 Evaluating Existing Explanations

The pieces of evidence obtained from the previous section can generally be summarized as follows:

- Relative volatility primarily increased in services industries, particularly trade, FIRE+ and TSC
- Relative volatility changes have been occurring throughout time, with large jumps in the 1980s, but also recent declines (trade, TSC) and increases (FIRE+, construction, mining, utilities) throughout the past 30 years.
- Changes in the relative co-movements of between industry hours and between industry outputs have also played an important role.

I now explore how some of the proposed changes in labor frictions responsible for the acyclicality of productivity may or may not be consistent with this evidence.

A commonly proposed explanation is the decline in unionization rates, particularly the sharp decline of union formation occurring around the 1980s. Hirsch (2008) document that unionization is greatest in industries such as manufacturing, construction, utilities and transportation and communication services, but that there was a decline in the share of unionized employment both in these union-heavy industries and in the remaining industries. Given that the decline is fairly widespread, one might expect that relative volatility would increase in these industries where there is the greatest density of unionized employment. However, of the industries discussed, only transportation and communication services saw a significant increase in relative volatility in the 1980s. Construction and utilities saw some gradual or eventual increase, but manufacturing even less.

Perhaps the unionization explanation can explain changes in relative volatility in utilities and manufacturing, but the timing doesn't completely coincide in construction, where much of the decline of unionization was concentrated in the 1980s, nor does the unionization explanation seem consistent with the sudden reversal in relative volatility in the transportation/storage/communication services industry. This is even without mentioning the large increases in relative volatility occurring in trade and FIRE+, which are far less unionized. Thus, the unionization story sees unlikely to explain some of the central features of acyclical productivity.

On the other hand, explanations regarding the rise of pay for performance and online recruitment practices seem more apt to be consistent with the increases in relative volatility in the services industries. Lemieux et al. (2009), for example, documents that pay for performance practices tend to be more prevalent in these services industries. Whether these industries are where pay for performance practices have most increased lately, is unclear however, as there is not immediately available evidence to this point. These two recent trends seem quite plausible for explaining the recent continued rise of relative volatility in the FIRE+ industry and the perhaps also the gradual rises in mining, construction and utilities, though seem less feasible for the temporary increases observed in TSC or trade. And at least as the timing goes, online recruitment practices certainly can't explain the large jump in the 1980s, as these were not heavily used until the mid-1990s.

Generally speaking, labor friction-centric stories appear to have a hard time explaining all the dimensions of change in relative volatility at the industry level. And to this point, there has been no added discussion in the changes in relative between industry correlations, which additionally seem more difficult to reconcile with stories of labor frictions. While there are elements of the recent acyclicality productivity which could be explained by these changes, it seems like there is yet work to be done in identifying how these could directly relate to the observed changes and especially yet work to be done in explaining the changes in relative industry co-movements of hours and output.

3.6 Conclusions

In this chapter, I have made several points regarding the recent acyclicality of labor productivity which have important implications for existing explanations of this phenomenon. In particular, acyclical productivity is synonymous with the recent change in the relative volatility of hours and output. Further, using industry level data, there appears to be substantial heterogeneity in changes in industry level relative volatilities across industries and across time and an important role for changes in the co-movements between industries. These pieces of evidence appear to be difficult to wholly reconcile with stories regarding labor frictions, which have been the most common story to emerge in explaining acyclical labor productivity thus far.

So what kind of stories seem to be more promising? Given the close connection between the Great Moderation and acyclical labor productivity, it seems like an important piece to understanding the variations in relative volatility across industries, and thus aggregate acyclical productivity, seems to be understanding the changing composition of shocks in the economy in the Great Moderation. The changes in volatility occurring in the Great Moderation have been concentrated in some industries more than others, and it is possible that a better understanding of how these shocks have changed in individual industries have propagated to the rest of the economy appears important, especially given the observed changes in industry output co-movements. Research focusing on industry asymmetries in the Great Moderation, the changing composition of shocks and changes in the propagation channels and co-movement structure of the economy seem particularly important in furthering our understanding of acyclical productivity. Recent examples along these lines include Foerster et al. (2011), Acemoglu et al. (2012) and Garin et al. (2013).

To this point, however, the clear conclusion of this chapter is that when industry and time variation is accounted for, the puzzle that is acyclical productivity is even more complicated than it appeared before.

3.7 Appendix

3.7.1 Proof of Lemma Relating Relative Volatility and Output-Productivity Correlation

Begin with the definition of correlation for output and labor productivity, where σ_x denotes the standard deviation of variable x:

$$\rho_{y,y-h} \equiv \frac{COV(y,y-h)}{\sigma_y \sigma_{y-h}}$$

Using the linearity properties of covariance and rearranging terms, we can write:

$$\frac{COV(y,y-h)}{\sigma_y \sigma_{y-h}} = \frac{COV(y,y) - COV(y,h)}{\sigma_y \sigma_{y-h}} = \frac{\sigma_y^2}{\sigma_y \sigma_{y-h}} - \frac{COV(y,h)}{\sigma_y \sigma_{y-h}} = \frac{\sigma_y}{\sigma_{y-h}} - \frac{\sigma_h}{\sigma_{y-h}}\rho_{y,h} = \frac{\sigma_y}{\sigma_{y-h}}\left(1 - \frac{\sigma_h}{\sigma_y}\rho_{y,h}\right)$$

Then, given the linear properties of variance, we know that we can write σ_{y-h} as:

$$\sigma_{y-h} = \sqrt{\sigma_y^2 + \sigma_h^2 - 2COV(y,h)} = \sqrt{\sigma_y^2 + \sigma_h^2 - 2\sigma_y \sigma_h \rho_{y,h}} = \sigma_y \sqrt{1 + \left(\frac{\sigma_h}{\sigma_y}\right)^2 - 2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}}$$

Combining this with the above expression, we get that:

$$\frac{\sigma_y}{\sigma_{y-h}}\left(1 - \frac{\sigma_h}{\sigma_y}\rho_{y,h}\right) = \frac{1 - \frac{\sigma_h}{\sigma_y}\rho_{y,h}}{\sqrt{1 + \left(\frac{\sigma_h}{\sigma_y}\right)^2 - 2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}}} \equiv f(\frac{\sigma_h}{\sigma_y}, \rho_{y,h})$$

Thus, the correlation of output and productivity can be written as purely a function of the

relative volatility of output and hours and the correlation between the two[10].

Showing that this functions is decreasing in $\frac{\sigma_h}{\sigma_y}$ is straightforward. Taking the partial derivative with respect to $\frac{\sigma_h}{\sigma_y}$, we get the following expressions:

$$\begin{aligned}
\frac{\partial f}{\partial \frac{\sigma_h}{\sigma_y}} &= -\frac{\rho_{y,h}}{\sqrt{1+\left(\frac{\sigma_h}{\sigma_y}\right)^2-2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}}}+\frac{\left(1-\frac{\sigma_h}{\sigma_y}\rho_{y,h}\right)\left(\rho_{y,h}-\frac{\sigma_h}{\sigma_y}\right)}{\left(1+\left(\frac{\sigma_h}{\sigma_y}\right)^2-2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}\right)^{3/2}} \\
&= \frac{-\rho_{y,h}\left(1+\left(\frac{\sigma_h}{\sigma_y}\right)^2-2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}\right)+\left(1-\frac{\sigma_h}{\sigma_y}\rho_{y,h}\right)\left(\rho_{y,h}-\frac{\sigma_h}{\sigma_y}\right)}{\left(1+\left(\frac{\sigma_h}{\sigma_y}\right)^2-2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}\right)^{3/2}} \\
&= \frac{\frac{\sigma_h}{\sigma_y}\left(\rho_{y,h}^2-1\right)}{\left(1+\left(\frac{\sigma_h}{\sigma_y}\right)^2-2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}\right)^{3/2}}
\end{aligned}$$

Given that the denominator is positive by definition of the standard deviation, as is the relative volatility in the numerator, since $|\rho_{y,h}| \leq 1$ by definition of the correlation coefficient, f is weakly decreasing in $\frac{\sigma_h}{\sigma_y}$. Thus completes the proof. □

[10] The above operations can similarly be done for the co-movement between hours and productivity, obtaining the following expression:

$$\rho_{h,y-h} = \frac{\rho_{y,h}-\frac{\sigma_h}{\sigma_y}}{\sqrt{1+\left(\frac{\sigma_h}{\sigma_y}\right)^2-2\left(\frac{\sigma_h}{\sigma_y}\right)\rho_{y,h}}} \equiv g(\frac{\sigma_h}{\sigma_y},\rho_{y,h})$$

Generally speaking, the co-movement between any object m and labor productivity can be written as a function $g(\rho_{y,h},\frac{\sigma_h}{\sigma_y},\frac{\rho_{y,m}}{\rho_{h,m}})$.

3.7.2 Complete Two-Sector Decomposition of Aggregate Relative Volatility

Because of the complexity of the full decomposition of aggregate relative volatility, I consider just two industries - goods and services. I define the goods sector to be given by the sum of the manufacturing, construction and mining sectors; the services sector is defined as the sum of the remaining five sectors. The relative volatility in terms of just these two industries is given by:[11]

$$\frac{Var\left(\widehat{ln(H)}\right)}{Var\left(\widehat{ln(Y)}\right)} = \frac{\sigma_H^2}{\sigma_Y^2} = \frac{\left(\frac{h_g}{H}\right)^2 \sigma_{H,g}^2 + \left(\frac{h_s}{H}\right)^2 \sigma_{H,s}^2 + 2\left(\frac{h_g}{H}\right)\left(\frac{h_s}{H}\right) Cov\left(\widehat{ln(h_g)}, \widehat{ln(h_s)}\right)}{\left(\frac{y_g}{Y}\right)^2 \sigma_{Y,g}^2 + \left(\frac{y_s}{Y}\right)^2 \sigma_{Y,s}^2 + 2\left(\frac{y_g}{Y}\right)\left(\frac{y_s}{Y}\right) Cov\left(\widehat{ln(y_g)}, \widehat{ln(y_s)}\right)}$$

As we are interested in the aggregate relative volatility, it is desirable to be able to rewrite the above expression in terms of relative volatilities and relative covariances at the industry level. With some simple algebra, one can rewrite the above as follows:

$$\frac{\sigma_H^2}{\sigma_Y^2} = \frac{\left(\left(\frac{y_g}{Y}\right)\sigma_{Y,g}\right)^2}{\sigma_Y^2}\left(\frac{\frac{h_g}{H}\sigma_{H,g}}{\frac{y_g}{Y}\sigma_{Y,g}}\right)^2 + \frac{\left(\left(\frac{y_s}{Y}\right)\sigma_{Y,s}\right)^2}{\sigma_Y^2}\left(\frac{\frac{h_s}{H}\sigma_{H,s}}{\frac{y_s}{Y}\sigma_{Y,s}}\right)^2 + \left(1 - \frac{\left(\left(\frac{y_g}{Y}\right)\sigma_{Y,g}\right)^2}{\sigma_Y^2} - \frac{\left(\left(\frac{y_s}{Y}\right)\sigma_{Y,s}\right)^2}{\sigma_Y^2}\right)\frac{\frac{h_g}{H}\frac{h_s}{H}}{\frac{y_g}{Y}\frac{y_s}{Y}}\frac{Cov\left(\widehat{ln(h_g)}, \widehat{ln(h_s)}\right)}{Cov\left(\widehat{ln(y_g)}, \widehat{ln(y_s)}\right)}$$

The above expression casts aggregate relative volatility as a weighted sum of weighted relative variances and weighted relative covariances. The outer weights on this sum represent the share of output volatility attributable to each industry and to the covariance between them. Changes in the output volatility weights are affected by the relative output volatilities of services and goods, their covariance, and their relative sizes. As the relative covariance terms reflect both relative variances and relative correlations, I rewrite the covariances in terms of variances and correlations, and get the following final expression for the decomposition:

[11] In the exercises that follow, I will technically allow for each window in which I compute the variance to be based on a different approximation point. However, this only affects the shares, and if anything, makes them more "flexible," generating an even better aggregate decomposition.

$$\frac{\sigma_H^2}{\sigma_Y^2} = \frac{\left(\left(\frac{y_g}{Y}\right)\sigma_{Y,g}\right)^2}{\sigma_Y^2}\left(\frac{\frac{h_g}{H}\sigma_{H,g}}{\frac{y_g}{Y}\sigma_{Y,g}}\right)^2 + \frac{\left(\left(\frac{y_s}{Y}\right)\sigma_{Y,s}\right)^2}{\sigma_Y^2}\left(\frac{\frac{h_s}{H}\sigma_{H,s}}{\frac{y_s}{Y}\sigma_{Y,s}}\right)^2 + \left(1 - \frac{\left(\left(\frac{y_g}{Y}\right)\sigma_{Y,g}\right)^2}{\sigma_Y^2} - \frac{\left(\left(\frac{y_s}{Y}\right)\sigma_{Y,s}\right)^2}{\sigma_Y^2}\right)\left(\frac{\frac{h_g}{H}\sigma_{H,g}}{\frac{y_g}{Y}\sigma_{Y,g}}\right)\left(\frac{\frac{h_s}{H}\sigma_{H,s}}{\frac{y_s}{Y}\sigma_{Y,s}}\right)\frac{\rho_{g,s}^h}{\rho_{g,s}^y}$$

Thus, the final decomposition of aggregate relative volatility reflects changes in: within industry relative volatilities, between industry relative correlations, relative hours and output shares of the economy, and the output volatility weights.

Table 3.7 reports the values for the sector shares of hours and output, sector volatilities in hours and output, and between sector correlations for the periods 1953-1982 and 1983-2004 and Table 3.8 then presents the relative volatilities, relative shares, relative correlations and output volatility weights for these same time periods. From Table 3.8, we observe several well-known trends in these two sectors - the goods sector is declining as a share of the total economy while services is rising, output and hours volatility declines in goods were greater than in services, and the correlation between goods and services outputs has fallen some in the Great Moderation. Table 3.8 then presents these changes in the relative terms which appear in the above decomposition. In terms of relative changes, the largest changes have come in the relative volatility of hours to output for the services industry and the relative correlation of hours and output between goods and services - both have risen by 58% in the latter period. Looking back at Table 3.3, we see that this is because while the volatility of services output has fallen by 25%, its hours volatility has risen by 19%. Similarly, the correlation between outputs for goods and services has fallen by a third, whereas the correlation between hours in these two industries has slightly risen. In contrast, the relative volatility of hours and output in the goods industry has barely changed at all, increasing by a mere 7%, and there have been virtually no changes in the relative hours and output shares attributable to either goods or services. The final component that we observe in the decomposition, the output volatility weights, requires a little more explanation. The output volatility weight for goods has declined quite a bit while the output volatility weight for services

Table 3.7: Descriptive Statistics on Industry Shares, Volatilities and Correlations for Goods and Services

	1953-1982	1983-2004	%Δ
Hours Share			
Goods	0.428	0.327	-24%
Serv	0.572	0.674	18%
Output Share			
Goods	0.379	0.291	-23%
Serv	0.621	0.709	14%
Hours Volatility			
Goods	0.031	0.020	-40%
Serv	0.010	0.012	-25%
Output Volatility			
Goods	0.037	0.022	-36%
Serv	0.013	0.010	19%
Hours Correlation	0.895	0.945	6%
Output Correlation	0.911	0.608	-33%

has risen substantially - this is largely due to the structural transformation from goods to services that has been gradually occurring over the past 50-60 years. But this change is also due in part to the fact that the volatility of output in goods has declined more than the volatility of output in services. The decline in weight on the covariance term occurs for similar reasons - the covariance between goods and services output has fallen by more than the variance of output in the services sector (and also the goods sector as well).

Just from looking at these tables, the most significant changes for aggregate relative volatility seem to be changes in the relative volatility in the services sector, as well as changes in the relative correlations of hours and output for the goods and services sectors, with perhaps a role for the changing composition of output volatility. Placing an exact percentage on how much each of these changes has contributed to the overall change is slightly trickier, however. One simple approach is to consider how much of the aggregate relative volatility change would have been observed if only one of the components had changed - essentially a counterfactual decomposition.

A common question in economics is how to decompose changes of a function into changes in the inputs of a function. When the function and its components are additively separable (for

Table 3.8: Relative Shares, Volatilities and Correlation for Goods and Services

	1953-1982	1983-2004	%Δ
Relative Shares			
Goods	1.129	1.124	-1%
Serv	0.921	0.951	3%
Relative Volatility			
Goods	0.834	0.890	7%
Serv	0.746	1.178	58%
Relative Corr.	0.982	1.554	58%
Output Vol. Weight			
Goods	0.409	0.276	-33%
Serv	0.146	0.348	138%
Cov	0.445	0.377	-15%

example, in the growth accounting case), this exercise is relatively straightforward. Unfortunately, there can arise a variety of settings in which the function of interest is not additively separable in its inputs. A common approach to take in such settings is to consider some form of counterfactual decomposition. A simple way of thinking about this problem is to consider attributing changes in a non-additively separable function $f(x,y,z)$ to changes in x,y, and z . Generally speaking, one way to decompose these changes is the following:

$$f(x',y',z') - f(x,y,z) =$$
$$\underbrace{f(x',y',z') - f(x',y',z)}_{\text{"}\Delta z\text{"}} + \underbrace{f(x',y',z) - f(x',y,z)}_{\text{"}\Delta y\text{"}} + \underbrace{f(x',y,z) - f(x,y,z)}_{\text{"}\Delta x\text{"}}$$

The idea is to hold x or y fixed at either its initial or terminal value and then evaluate the change in the function f that occurs from shifting the other variable. This is essentially the spirit of the Oaxaca-Blinder decomposition for wage changes, as well as the spirit of identifying a price index using either the Laspeyres or Paasche method. It is also commonly called structural decomposition analysis (SDA) in the energy economics and input-output literature. The challenge with these

methods is that they are sensitive to the reference point of the decomposition, that is, the values you hold the inputs fixed at while evaluating the change. And there are a variety of possible reference points, meaning for a function with n inputs, there are $n!$ different possible decompositions.

One approach to compensate for this problem is simply to calculate the average of the contribution of a given input across holding all other values at the terminal values or holding all other values at their initial values. Dietzenbacher and Los (1998) show that this is often a reasonable approximation, with minimal loss relative to computing all the possible decompositions. In the above decomposition, this implies[12]:

$$\begin{aligned}
\text{"}\Delta z\text{"} &= 0.5(f(x',y',z') - f(x',y',z)) + 0.5(f(x,y,z') - f(x,y,z)) \\
\text{"}\Delta y\text{"} &= 0.5(f(x',y',z') - f(x',y,z')) + 0.5(f(x,y',z) - f(x,y,z)) \\
\text{"}\Delta x\text{"} &= 0.5(f(x',y',z') - f(x',y,z')) + 0.5(f(x,y',z) - f(x,y,z))
\end{aligned}$$

For this decomposition, I consider changes in aggregate volatility driven by changes in six separate components: each of the two relative volatilities, the relative correlation, changes in the relative sector shares, changes in the output volatility weights stemming from changes in sector sizes and changes in the same weights stemming from changes in output variances and covariance[13]. The results of this decomposition are reported in Table 3.9. As expected, the changes in relative volatility in services and the change in relative correlations are most important for the change in aggregate relative volatility, explaining roughly 90% of the change in aggregate relative volatility. Goods relative volatility only contributes a small amount to the aggregate change, just over 10%. Given the somewhat large changes in the output volatility weights, it is perhaps surprising that these are not particularly important for the aggregate change. The reasoning for this result

[12]It is worth noting that this approximation does not actually constrain the sum of all changes to equal 100%. The closer the sum is to 100%, however, the better the approximation is apt to be.

[13]These output volatility composition changes are considered independently of changes in the relative volatilities or relative correlations

Table 3.9: Contributions of Two Sector Decomposition Components to Aggregate Change in Relative Volatility

	Contribution to Aggregate Relative Volatility Change
Goods Relative Volatility	11%
Serv Relative Volatility	59%
Relative Correlation	31%
Relative Shares	3.5%
Volatility Weight - Size	-1%
Volatility Weight - Volatility	-1%

is the following. If one were only to change the output volatility weights, this places much more weight on volatility in the services sector. However, prior to the Great Moderation, the relative volatilities in the services sector were lower than in the goods sector, meaning that a pure compositional change will actually lead to aggregate decreases in relative volatility. And on the other hand, had all the components of the decomposition changed except the output volatility weights, that would have placed much greater weight on the changes in the goods sector and additional weight on the substantial changes in the covariance term, largely compensating for the reduced weight on services. Thus, while there is potentially some role for these changes in volatility composition, it is not clear how important they are, since independent of other changes, they do not lead to aggregate increases in relative volatility.

Bibliography

ACEMOGLU, D. AND D. H. AUTOR (2011): "Skills, Tasks and Technologies: Implications for Employment and Earnings," *Handbook of Labor Economics*, 4, 1043–1171.

ACEMOGLU, D., V. M. CARVALHO, A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80, 1977–2016.

ATKINSON, A. B., T. PIKETTY, AND E. SAEZ (2011): "Top Incomes in the Long Run of History," *Journal of Economic Literature*, 49, 3–71.

ATTANASIO, O., H. LOW, AND V. SANCHEZ-MARCOS (2008): "Explaining Changes in Female Labor Supply in a Life-cycle Model," *American Economic Review*, 98, 1517–1552.

AUTOR, D. H. (2010): "The Polarization of Job Opportunities in the US Labor Market: Implications for Employment and Earnings," *Center for American Progress and The Hamilton Project*.

AUTOR, D. H. AND D. DORN (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103, 1553–1597.

AUTOR, D. H., D. DORN, AND G. H. HANSON (2013a): "The China Syndrome: Local Labor Market Effects of Import Competition in the United States," *American Economic Review*, 103, 2121–2168.

——— (2013b): "Untangling Trade and Technology: Evidence from Local Labor Markets," *NBER Working Paper No. 18938*.

AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): "The Polarization of the US Labor Market," *American Economic Review: Papers and Proceedings*, 96, 189–194.

——— (2008): "Trends in US Wage Inequality: Revising the Revisionists," *Review of Economics and Statistics*, 90, 300–323.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118, 1279–1333.

BARNICHON, R. (2010): "Productivity and Unemployment Over the Business Cycle," *Journal of Monetary Economics*, 57, 1013–1025.

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

BEAUDRY, P., D. A. GREEN, AND B. M. SAND (2013): "The Great Reversal in the Demand for Skill and Cognitive Tasks," NBER Working Paper No. 18901.

BERGER, D. (2012): "Countercyclical Restructuring and Jobless Recoveries," Working Paper.

BOSKIN, M. J., E. R. DULBERGER, Z. GRILICHES, R. J. GORDON, AND D. JORGENSEN (1996): *Toward a More Accurate Measure of the Cost of Living: Final Report to the Senate Finance Committee for the Advisory Commision to Study the Consumer Price Index.*

CORTES, G. M. (2012): "Where Have the Routine Workers Gone? A Study of Polarization using Panel Data," Working Paper.

CORTES, G. M., N. JAIMOVICH, C. NEKARDA, AND H. E. SIU (2012): "The Micro and Macro of Job Polarization," Working Paper.

CUMMINS, J. G. AND G. L. VIOLANTE (2002): "Investment-specific Technical Change in the United States (1947–2000): Measurement and Macroeconomic Consequences," *Review of Economic Dynamics*, 5, 243–284.

DICECIO, R. (2009): "Sticky Wages and Sectoral Labor Comovement," *Journal of Economic Dynamics and Control*, 33, 538–553.

DIETZENBACHER, E. AND B. LOS (1998): "Structural Decomposition Techniques: Sense and Sensitivity," *Economic Systems Research*, 10, 307–324.

DORN, D. (2009): "Essays on Inequality, Spatial Interaction, and the Demand for Skills," Ph.D. thesis, University of St. Gallen.

ELSBY, M. W., B. HOBIJN, AND A. SAHIN (2013): "The Decline of the US Labor Share," *Brookings Papers on Economic Activity*.

ESTRADA, A. AND E. VALDEOLIVAS (2012): "The Fall of the Labour Income Share in Advanced Economies," *Occasional Paper Series 1209*, Bank of Spain.

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, 1–38.

FOGLI, A. AND L. VELDKAMP (2011): "Nature or Nurture? Learning and the Geography of Female Labor Force Participation," *Econometrica*, 79, 1103–1138.

FOOTE, C. AND R. W. RYAN (2012): "Labor-market Polarization Over the Business Cycle," *Federal Reserve Bank of Boston Public Policy Discussion Series*.

FRYDMAN, C. AND D. JENTER (2010): "CEO Compensation," *Annual Review of Financial Economics*, 2, 75–102.

GALÍ, J. (2011): "The Return of the Wage Phillips Curve," *Journal of the European Economic Association*, 9, 436–461.

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

GALÍ, J. AND T. VAN RENS (2010): "The Vanishing Procyclicality of Labor Productivity," Working Paper.

GARIN, J., M. PRIES, AND E. SIMS (2013): "Reallocation and the Changing Nature of Economic Fluctuations," Working Paper.

GOOS, M. AND A. MANNING (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *Review of Economics and Statistics*, 89, 118–133.

GOOS, M., A. MANNING, AND A. SALOMONS (2011): "Explaining Job Polarization: The Roles of Technology, Offshoring and Institutions," *CES-Discussion Paper Series DPS11. 34*, 1–35.

GORDON, R. J. (1990): *The Measurement of Durable Goods Prices*, University of Chicago Press.

——— (2006): "The Boskin Commission Report: A Retrospective One Decade Later," *International Productivity Monitor*, 12, 7–22.

GREENWOOD, J., Z. HERCOWITZ, AND P. KRUSELL (1997): "Long-Run Implications of Investment-Specific Technological Change," *American Economic Review*, 87, 342–362.

GROSSMAN, G. M. AND E. ROSSI-HANSBERG (2008): "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98, 1978–1997.

HAGEDORN, M. AND I. MANOVSKII (2011): "Productivity and the Labor Market: Comovement over the Business Cycle," *International Economic Review*, 52, 603–619.

HAMERMESH, D. S. (1996): *Labor Demand*, Princeton University Press.

HANSEN, G. D. AND R. WRIGHT (1992): "The Labor Market in Real Business Cycle Theory," *Federal Reserve Bank of Minneapolis Quarterly Review*, 16, 2–12.

HE, H. AND Z. LIU (2008): "Investment-specific Technological Change, Skill Accumulation, and Wage Inequality," *Review of Economic Dynamics*, 11, 314–334.

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, 15–51.

HIRSCH, B. T. (2008): "Sluggish Institutions in a Dynamic World: Can Unions and Industrial Competition Coexist?" *Journal of Economic Perspectives*, 22, 153–176.

HURST, E., G. LI, AND B. PUGSLEY (2014): "Are Household Surveys Like Tax Forms? Evidence from Income Underreporting of the Self-Employed," *Review of Economics and Statistics*, 96, 19–33.

JAIMOVICH, N. AND H. E. SIU (2012): "The Trend is the Cycle: Job Polarization and Jobless Recoveries," *NBER Working Paper 18334*.

JUHN, C. AND S. POTTER (2006): "Changes in Labor Force Participation in the United States," *Journal of Economic Perspectives*, 20, 27–46.

JUNG, J. AND J. MERCENIER (2010): "Routinization-biased Technical Change, Globalization and Labor Market Polarization: Does Theory fit the Facts?" *Tepp Working Paper No. 2011-10*.

KALDOR, N. (1961): *Capital Accumulation and Economic Growth*, Macmillan.

KARABARBOUNIS, L. AND B. NEIMAN (2014): "The Global Decline of the Labor Share," *Quarterly Journal of Economics*, 129, 61–103.

KATZ, L. F. AND H. DAVID (1999): "Changes in the Wage Structure and Earnings Inequality," *Handbook of Labor Economics*, 3, 1463–1555.

KING, M., S. RUGGLES, J. T. ALEXANDER, S. FLOOD, K. GENADEK, M. B. SCHROEDER, B. TRAMPE, AND R. VICK (2010): "Integrated Public Use Microdata Series, Current Population Survey: Version 3.0," [Machine-readable database]. Minneapolis: University of Minnesota.

KRUSELL, P., L. E. OHANIAN, J.-V. RÍOS-RULL, AND G. L. VIOLANTE (2000): "Capital-skill Complementarity and Inequality: A Macroeconomic Analysis," *Econometrica*, 68, 1029–1053.

LAMBERT, R. A. AND D. F. LARCKER (2001): "How Employees Value (Often Incorrectly) Their Stock Options," Knowledge@Wharton Online Article, available at http://knowledge.wharton.upenn.edu.

LARRIMORE, J., R. V. BURKHAUSER, S. FENG, AND L. ZAYATZ (2008): "Consistent Cell Means for Topcoded Incomes in the Public Use March CPS (1976–2007)," *Journal of Economic and Social Measurement*, 33, 89–128.

LEMIEUX, T., W. B. MACLEOD, AND D. PARENT (2009): "Performance Pay and Wage Inequality," *Quarterly Journal of Economics*, 124, 1–49.

LIPTON, D., J. POTERBA, J. SACHS, AND L. SUMMERS (1982): "Multiple Shooting in Rational Expectations Models," *Econometrica*, 1329–1333.

MALIAR, L. AND S. MALIAR (2011): "Capital–Skill Complementarity and Balanced Growth," *Economica*, 78, 240–259.

MCCULLY, C. P. (2013): "Integration of Micro and Macro Data on Consumer Income and Expenditures," in *Measuring Economic Sustainability and Progress*.

MEYER, B. D. AND D. T. ROSENBAUM (2001): "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers," *Quarterly Journal of Economics*, 116, 1063–1114.

MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2014): "Has ICT Polarized Skill Demand?: Evidence from Eleven Countries Over 25 Years," *Review of Economics and Statistics*, 96, 60–77.

MISHEL, L. AND N. SABADISH (2013): "Methodology for Measuring CEO Compensation and the Ratio of CEO-to-worker Compensation," *Economic Policy Institute Working Paper*, 298.

MOFFITT, R. A. (2012): "The US Employment-Population Reversal in the 2000s: Facts and Explanations," *NBER Working Paper No. 18520*.

MOYLAN, C. E. (2008): "Employee Stock Options and the National Economic Accounts," *Survey of Current Business*, 88, 7–13.

NUCCI, F. AND M. RIGGI (2013): "Performance Pay and Changes in US Labor Market Dynamics," *Journal of Economic Dynamics and Control*, 37, 2796–2813.

RAVN, M. O. AND H. UHLIG (2002): "On Adjusting the Hodrick-Prescott filter for the Frequency of Observations," *Review of Economics and Statistics*, 84, 371–376.

SMITH, C. L. (2013): "The Dynamics of Labor Market Polarization," Tech. rep., Board of Governors of the Federal Reserve System (US).

STIROH, K. J. (2009): "Volatility Accounting: A Production Perspective on Increased Economic Stability," *Journal of the European Economic Association*, 7, 671–696.

STOCK, J. H. AND M. W. WATSON (2003): "Has the Business Cycle Changed and Why?" in *NBER Macroeconomics Annual 2002, Volume 17*, MIT press, 159–230.

WARNOCK, M. C. AND F. E. WARNOCK (2000): "The Declining Volatility of US Employment: Was Arthur Burns Right?" 677.

YOUNG, A. T. (2010): "One of the Things We Know That Ain't So: Is US Labor's Share Relatively Stable?" *Journal of Macroeconomics*, 32, 90–102.