

Essays on Economic Growth and Structural Transformation

by

Alberto José Vindas Quesada

A Dissertation Presented in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Approved March 2019 by the
Graduate Supervisory Committee:

Bart Hobijn, Chair
Alexander Bick
Gustavo Ventura

ARIZONA STATE UNIVERSITY

May 2019

ABSTRACT

This dissertation consists of three essays on modern economic growth and structural transformation, in particular touching on the reallocation of labor across industries, occupations, and employment statuses.

The first chapter investigates the quantitative importance of non-employment in the labor market outcomes for the United States. During the last 50 years, production has shifted from goods to services. In terms of occupations, the routine employment share decreased, giving way to increases in manual and abstract ones. These two patterns are related, and lower non-employment had an important role. A labor allocation model where goods, market services, and home services use different tasks as inputs is used for quantitative exercises. These show that non-employment could significantly slow down polarization and structural transformation, and induce significant displacement within the labor force.

The second chapter, coauthored with Bart Hobijn and Todd Schoellman, looks at the demographic structure of structural transformation. More than half of labor reallocation during structural transformation is due to new cohorts disproportionately entering growing industries. This suggests substantial costs to labor reallocation. A model of overlapping generations with life-cycle career choice under switching costs and structural transformation is studied. Switching costs accelerate structural transformation, since forward-looking workers enter growing industries in anticipation of future wage growth. Most of the impact of switching costs shows on relative wages.

The third chapter establishes that job polarization is a global phenomenon. The analysis of polarization is extended from a group of developed countries to a sample of 119 economies. At all levels of development, employment shares in routine occupations have decreased since the 1980s. This suggests that routine occupations are becoming increasingly obsolete throughout the world, rather than being outsourced

to developing countries. A development accounting framework with technical change at the *task* level is proposed. This allows to quantify and extrapolate task-specific productivity levels. Recent technological change is biased against routine occupations and in favor of manual occupations. This implies that in the following decades, world polarization will continue: employment in routine occupations will decrease, and the reallocation will happen mostly from routine to manual occupations, rather than to abstract ones.

DEDICATION

To my grandfather, Jorge.

ACKNOWLEDGMENTS

This dissertation was possible as a result of my interactions with several people in the economics department at ASU. I am deeply grateful to my advisor Bart Hobijn for the continuous motivation and advice, the pleasant and not-so-pleasant conversations, and most importantly, for his unparalleled commitment to my formation as an economist. Alex Bick provided sharp insights and questions throughout this process, which shaped my way of tackling these problems. Gustavo Ventura gave useful suggestions, kind advice, and valuable perspectives during the development of these essays.

I am greatly indebted to my fellow students in the department and faculty members for their support, questions, and discussion during the last five years.

TABLE OF CONTENTS

	Page
LIST OF TABLES	viii
LIST OF FIGURES	x
1 JOB POLARIZATION, STRUCTURAL TRANSFORMATION, AND NON-EMPLOYMENT	1
1.1 Introduction	1
1.2 Job Polarization and Non-employment Changes	4
1.2.1 Hollowing Out the Employment Distribution	5
1.2.2 Within Industry Job Polarization	7
1.2.3 Lower Non-employment Filled the Edges	12
1.3 A Labor Allocation Model	15
1.3.1 Preferences & Technology	16
1.3.2 Competitive Equilibrium Outcomes	18
1.4 Quantitative Results: How Non-employment Affects Polarization ...	28
1.4.1 Calibrating the Model	28
1.4.2 Assessing the Importance of Non-employment	31
1.5 Conclusions	37
2 STRUCTURAL TRANSFORMATION BY COHORT	38
2.1 Introduction	38
2.2 Cohort Effects and Structural Transformation	43
2.3 Structural Transformation with Career Choices	48
2.3.1 Households	49
2.3.2 Firms	51
2.3.3 Career Decisions and the Labor Supply	52
2.3.4 Equilibrium	55

CHAPTER	Page
2.4 The Impact of (Re-)Training Costs	58
2.4.1 Flexible Benchmark and Solution Method	59
2.4.2 Retraining Costs Accelerate Structural Transformation	64
2.4.3 Retraining Costs Increase with Age	66
2.4.4 Bulk of Impact of Retraining Costs Through Relative Wages	67
2.4.5 Output Losses Due to Retraining and Allocation of Labor ..	68
2.4.6 Unanticipated Acceleration of Structural Transformation....	70
2.5 Conclusion	72
3 WORLD POLARIZATION.....	81
3.1 Introduction.....	81
3.2 World Polarization Through the Lens of Occupational Development Profiles	84
3.3 Further Polarization Is on Its Way	92
3.3.1 Polarization Accounting	92
3.3.2 Accounting Results	96
3.3.3 VAR Analysis	100
3.3.4 Forecasting Results	106
3.4 Conclusions	116
REFERENCES	118
APPENDIX	
A APPENDICES TO CHAPTER 1	126
A.1 Data Sources	127
A.2 Detailed Tables.....	127
A.3 Model	129

APPENDIX	Page
A.3.1 Environment	129
A.3.2 Decentralized Market Structure	131
A.4 Estimation Procedure	132
B APPENDICES TO CHAPTER 2	134
B.1 United States Data	135
B.2 International Data	135
B.3 Alternative Industry Decompositions	136
B.4 Solution to Career Choices	138
B.5 Calibration of Flexible Benchmark Parameters	142
B.6 Details about Flexible Benchmark	143
B.7 Implementation of Extended-path Method	146
C APPENDICES TO CHAPTER 3	148
C.1 Data Sources	149
C.2 Shift-share Details	152
C.3 VAR Convergence System	152
C.4 A Model of Task Investment Specific Technical Change to Under- stand Polarization	153
C.4.1 Model Estimation	156
C.4.2 Alternative Identification Strategy	158

LIST OF TABLES

Table	Page
1.1 Employment Share Regression Results	7
1.2 Shift-Share Decomposition of Changes in Occupational Shares	12
1.3 Model Parameters	30
1.4 Productivity Estimates	31
1.5 Actual and Fitted Labor Shares	32
1.6 Freezing Non-employment, Counterfactual Results	34
1.7 Increasing Home Productivity, Counterfactual Results	35
2.1 Between Cohort Share of Structural Transformation	45
2.2 Between Demographic Group Share of Structural Transformation	48
2.3 ANOVA in Different Model Specifications	62
3.1 Occupational Development Profiles: 1985 & 2014	86
3.2 World Task-Specific Productivity Growth Rates (%): 1980-2014	97
3.3 U.S. Productivity Log-Differences: 1980-2014	102
3.4 VAR Estimate Results	105
3.5 U.S. Task-Specific Log-Productivity & Occupational Employment Shares (1980, 2014 & 2050)	107
3.6 World Productivity Gaps: 2014 & 2050	110
3.7 World Annual Occupational Employment Share Changes: 1980-2014 & 2014-2050	111
3.8 World Productivity Growth Sources	115
A.1 Occupational Job Polarization	127
A.2 Job Polarization and Non-employment	128
A.3 Occupation Shares within Industries	128
A.4 Industry Shares	128

Table	Page
B.1 Structural Transformation Sample.....	137
B.2 Between Cohort Share for Alternative Industry Groupings	137
C.1 Country Data Availability	150
C.2 Country Data Availability (Ctn'd)	151

LIST OF FIGURES

Figure	Page
1.1 Occupational Job Polarization	6
1.2 Occupation Shares Within Industries	9
1.3 Industry Shares	10
1.4 Occupational Job Polarization & Labor Non-Participation	14
2.1 Structural Transformation in the United States	40
2.2 Within cohort effects by age	46
2.3 Cohort-specific employment shares in three model specifications	75
2.4 Wages and Cohort career profiles in three model specifications	76
2.5 Sectoral employment shares in flexible benchmark and under flat re- training costs	77
2.6 Sectoral labor supply and labor demand curves at $t = 50$	78
2.7 Difference in output levels between increasing retraining costs and flex- ible benchmark	79
2.8 Impact of acceleration of structural transformation on employment and wages	80
3.1 Occupational Development Profiles: 1985 & 2014	87
3.2 Productivity Biases in 2014 with respect to 1985	98
3.3 U.S. Occupational Employment Shares: 1980 - 2050	108
3.4 World Occupational Employment Shares: 1985 - 2050	112
3.5 Occupational Development Profiles: 1985, 2014 & 2050	113
B.1 Countries in Dataset	136
B.2 Structural Transformation in the United States: 4 Industry View	138
B.1 Path of employment shares in flexible benchmark and data	144
B.1 Dynamics of flexible benchmark	145

Chapter 1

JOB POLARIZATION, STRUCTURAL TRANSFORMATION, AND NON-EMPLOYMENT

1.1 Introduction

The labor markets in the United States have changed significantly during the last 50 years. The share of routine occupations, heavy on procedural and repetitive tasks, decreased by 28%. On the other hand, manual occupations (heavy on physical tasks and in-person interactions), and abstract occupations (heavy on problem solving and creative tasks) increased by 27 and 45%. In terms of industries, the share of goods fell by 49%, giving way to an increase in services of 33%.¹ The first pattern is related to the *job polarization* process, the shrinking concentration of employment in routine occupations. The second pattern is related to the *structural transformation* process, the reallocation of economic activity across industries.

This transition has received considerable attention due to its implications on wage inequality and mobility costs. Routine occupations tend to be in the middle of the wage distribution, so job polarization pattern translates into higher wage differentials.² Human capital has been shown to have a heavy occupation-specific component, so displacements within the labor force would end up in costly adjustments for the workers.³

¹These are percentages with respect to their initial labor shares, so these don't add up to zero.

²Goos and Manning (2007) study this pattern in the United Kingdom, Autor and Dorn (2013) focus on the United States, while Goos *et al.* (2014) expand the analysis for 16 Western European countries.

³Kambourov and Manovskii (2009) document this in the United States using data from the Panel

This paper speaks to the latter point. Some of the job polarization analyses frame this transition in terms of job losses and disappearing routine occupations, which ensue large and costly reallocations.⁴ Using data from the Current Population Survey, I argue that this is not the case: the adjustment is mostly through decreases in non-employment. Between 1968 and 2018 non-employment decreased from 33 to 22%, a drop of 34%. This translated to an increase in abstract and manual occupations, which explains the decrease in routine's employment share.

This paper also speaks to the occupation-industry mix in the labor force. In quantitative terms, both the occupation mix within industries, and the industry mix within the economy play an important role in explaining overall polarization. Occupational changes within industries have a stronger effect on the increase in abstract occupations, while the shift towards services explains most of the increase in manual occupations.

In this article I ask about the quantitative importance of the decrease in non-employment for the productive structure of the economy. To answer that, I propose a labor allocation model explaining the occupational and industrial structure of the economy. It incorporates the non-employment decision, it justifies Baumol's cost disease from an occupational point of view, and gives the polarization process a treatment of the forces taking place between and within broad industries.

This model distinguishes between occupations in labor, and industries in consumption. Its building blocks are motivated by four patterns in the data. First, job polarization has played in a smooth, constant fashion during the last 50 years, so a persistent force should be behind these changes. Second, the adjustment implied a

Study of Income Dynamics.

⁴Examples include Autor (2010), Acemoglu and Autor (2011), Jaimovich and Siu (2012), and Mandelman (2016).

decrease in non-employment, which requires including this margin. Third, the occupational structure within goods and services differs substantially, so the industrial reallocation channel is of quantitative importance. Fourth, both the goods and services industries have polarized similarly, so the forces behind polarization should be occupation, rather than industry-specific.

To study the quantitative implications of the model, I calibrate it to the United States using data from 1968 to 2018. I find that productivity growth is the highest in routine occupations, followed by manual, abstract, and home production. The model is successful in reproducing the occupation dynamics within goods and market services, and is able to generate the movement towards market services we see in the data.

Finally, to assess the quantitative importance of lower non-employment, I perform two counterfactual exercises. The first exercise, inspired by women's insertion into the labor force, freezes non-employment at its 1968 level. This decreases the production of goods and market services by 2 and 18%, holds back structural transformation to its 1999 level, and decreases polarization by an average of 2%. The second exercise, inspired by the home productivity slowdown reported in Bridgman (2016), has home productivity growing at the rate of market services. This increases the production of goods by 17%, decreases the production of market services by 27%, holds back structural transformation to its 1977 level, and decreases polarization by an average of 7.5%.⁵ This illustrates not only the importance of this channel, but that its causes also play a important role.

The paper is organized as follows. Section 1.2 introduces the stylized facts behind job polarization. Section 1.3 presents the model I use. Section 1.4 deals with the quantitative matters: first it explains the estimation procedure, and then it goes over

⁵The comparison points are the model's predicted outcomes for 2018.

the counterfactual exercises. Finally, Section 3.4 concludes.

1.2 Job Polarization and Non-employment Changes

Between 1968 and 2018, the share of manual occupations in total employment increased by 27 percent, the share of routine occupations decreased by 28 percent, and the share of abstract occupations increased by 45 percent. This is to say, the occupational structure in the United States polarized. These large swings are described as worrisome due to possible displacements to lower paid occupations, as the share of routine occupations decreases, and a higher fraction of the employed population works in lower paying jobs (Jaimovich and Siu, 2012). Meanwhile, non-employment, the fraction of the population that is either unemployed or out of the labor force, decreased 34 percent. These are large changes, and this section documents them with more detail.

Accounting for non-employment's decrease suggests that the net flow of people from non-employment to employment account for polarization, not the flow of people from routine occupations to non-employment. Over this period, the share of the population working in routine occupations remained flat during most of these years, while the share in abstract and manual occupations increased. This is the novel, and main contribution of the empirical analysis: that polarization happened through decreases in non-employment. Few studies focus on this interaction. One notable exception is Cerina *et al.* (2017), that analyzes how female working hours shaped polarization. Even though the change in non-employment is one of the major adjustment margins, its effect has gone mostly unnoticed in the polarization literature.

I document these changes using data from the Annual Social and Economic (ASEC) supplement to the Current Population Survey. The focus is on the extensive margin of labor: whether people work in one of the three groups of occupations, or

whether they are non-employed. I consider the population aged between 25 and 65 years, and use their labels to determine industry, occupation (for the employed), and employment status. Lack of hourly data for all these years in CPS precludes studying the hours worked, and therefore, leisure. Fortunately, changes in leisure hours do not point at this being the leading force.⁶

In this paper I follow Cortés *et al.* (2014) and construct the occupational categories, grouped by their task content. Abstract occupations are intense in tasks that require problem solving, judgment, and creativity. Some examples are managers, lawyers, and architects. Routine occupations are heavy on tasks that follow precise, and well understood routinary procedures. Examples include cashiers, machinists and travel agents. Finally, manual occupations rely more on tasks that require flexibility, in-person interactions, and physical adaptability. These include janitors, bartenders, and nursing aides. Appendix A.1 presents a more detailed discussion of this data source and the classifications.


1.2.1 *Hollowing Out the Employment Distribution*

Job polarization has been described as a “hollowing out” of the occupational distribution, since routine occupations are also in the middle of the wage distribution. Indeed, the employment share of routine occupations decreased by 0.33 percentage points per year, which has happened steadily and as early as 1968.

The employment share of routine occupations fell 16.4 percentage points during the 1968-2018 period. Figure 1.1 shows this by plotting the employment shares for

⁶As Cerina *et al.* (2017) note, most of polarization can be attributed to women, but Aguiar and Hurst (2007), document that leisure has increased almost identically by gender. These studies using time-use surveys also note that total hours of work (including market and home work) are very similar across genders. This suggests that the intensive margin, at the aggregate level, is not the main force driving polarization, and as such is left out of this analysis.

Figure 1.1: Occupational Job Polarization



Shaded areas indicate years of major changes in occupational codes. These percentages refer to each occupation’s share in employment.

Source: author’s calculations using CPS.

these three occupations. The hollowing out is fairly evident: while manual occupations gained 3 percentage points, abstract occupations gained 13.4. By 2018, the employment shares of abstract and routine occupations were the same.

In addition, these changes have been smooth, and fairly constant over time. A regression analysis summarizes these points more clearly. Table 1.1 presents the results of regressing the employment shares against a constant term and a trend. All trends are statistically significant, show a good fit to the data with low dispersion around the regression lines. The constant trends mean that job polarization has been happening during the entire period.⁷ Several of the studies concerning polarization focus on more recent decades, and posit that it began in the 1980s. Since the CPS data

⁷The regression results for the first-differenced series throw very similar magnitudes, although the coefficient in manual occupations stops being statistically significant.

Table 1.1: Employment Share Regression Results

	Constant	Trend	R ²	Std. Error of Estimate
Manual	10.48	0.08	0.74	0.66
Routine	60.51	-0.33	0.96	1.05
Abstract	29.01	0.25	0.96	0.78

These are the results of regressing the occupation's employment share with a constant and a yearly trend term. All coefficients are statistically significant at 1%. The standard error of estimate is the standard deviation of the differences between the observed and predicted shares.

Source: author's calculations using CPS.

start in 1968, we can go further back in time and state that it has also been happening in earlier years. This is in line with the findings of Bárány and Siegel (2018), who use Census data to document the same pattern for an even longer period.

1.2.2 Within Industry Job Polarization

Structural transformation, the reallocation of employment across industries, can explain why polarization has been taking place. As industries that are more or less intensive in routine occupations grow, they can drive the occupational employment shares along with them. From an empirical point of view, can trends in this reallocation *between* industries explain most of polarization? In this section, I argue that the answer is no, that changes *within* these industries account for 66% of polarization. This means that most of polarization happens because each industry is demanding less routine workers, and more abstract workers. However, the changes between industries are also quantitatively important, in particular for manual occupations.

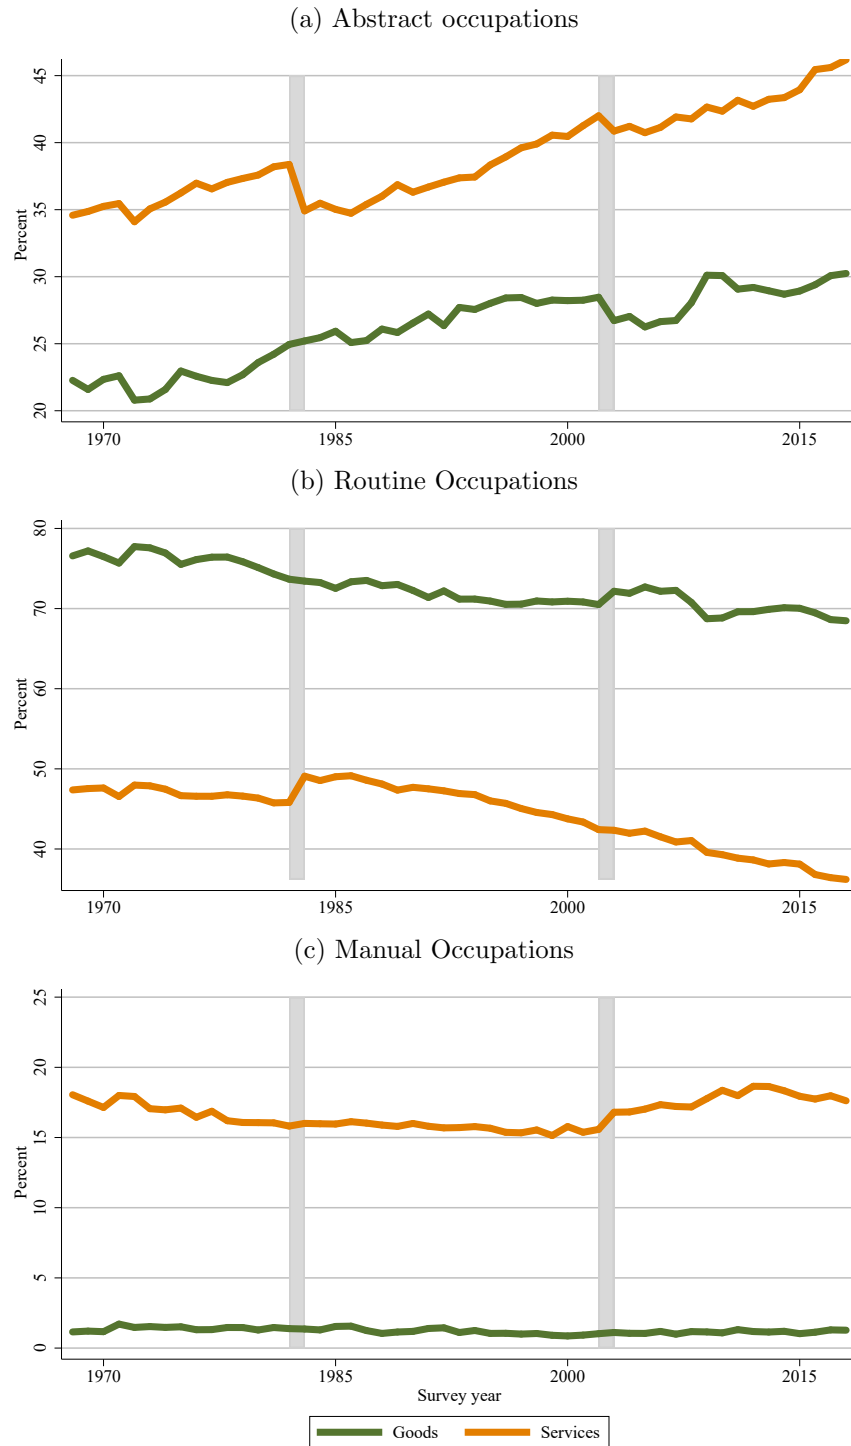
The focus of this section are occupational employment shares, but conditional on their industry: the goal is to establish the quantitative importance of the changes within each industry. I organize the productive structure into two industries: goods and services. Following the standard approach, the goods industry consists of agriculture and manufacturing, which include, among other categories, forestry and construction. The services industry includes categories such as retail, and professional and entertainment services.

Graphical examination of within industry occupational changes shows that these are very similar over time. Figure 1.2 plots the employment shares of each occupation within total employment in each industry. The first panel shows the shares of abstract occupations, which increase over time for both industries. The second shows the opposite effect for routine occupations, and the third shows that manual occupations remain relatively flat, especially in the goods industry. Although the changes were similar across industries, the magnitudes are larger for services: for the production of goods, the drop in routine occupations was 8.2 percentage points, while it was 11.2 in services. In abstract occupations, the production of goods increased by 8.1 percentage points, and 11.7 in services.

The changes within industries are in line with aggregate changes for routine and abstract occupations. For manual occupations, that is not the case: in the production of goods, its share increased 0.1 percentage points, and decreased 0.5 in the production of services, while overall, its employment share increased by 3. Within industry changes, then, cannot explain the increase in the overall employment share in manual occupations. To account for this, we need to turn to the changes in employment between industries.

Changes between industries can be quantitatively important if the employment shares within each industry are sufficiently different, and if we observe enough reallo-

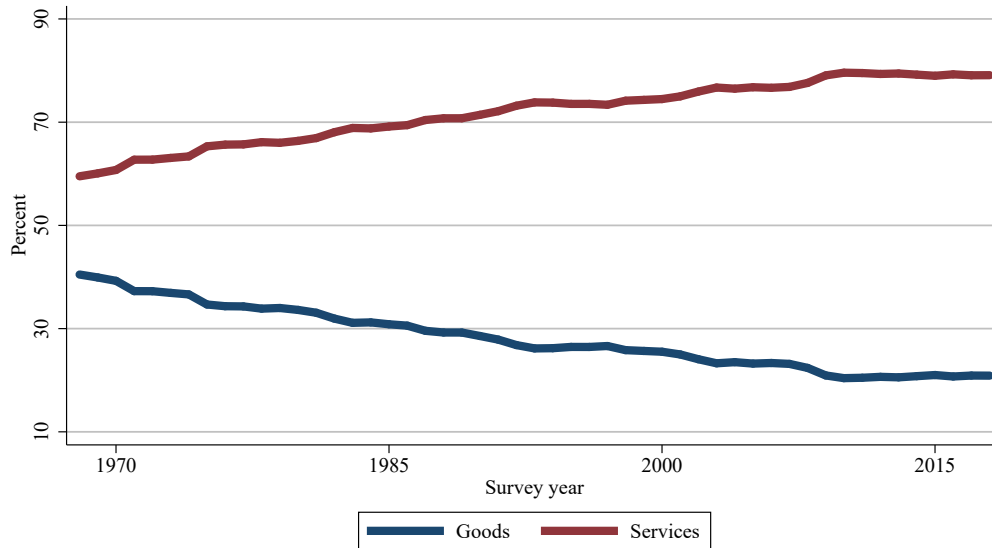
Figure 1.2: Occupation Shares Within Industries



Shaded areas indicate years of major changes in occupational codes. These percentages refer to the share of each occupation in each industry's total labor demand.

Source: author's calculations using CPS.

Figure 1.3: Industry Shares



These percentages refer to each industry's share in the labor force.

Source: author's calculations using CPS.

cation across industries. For manual occupations, both conditions are satisfied. The share of manual occupations in the production of services is 14 times as high as in goods, and the process of structural transformation increased the employment share in services by 19.7 percentage points. This last point is shown in Figure 1.3, that plots the industry employment shares over time.

To measure the relative importance of these movements, I perform a shift-share decomposition. Overall, changes within industries account for 66% of the occupational shifts, while between industry changes account for the remaining 34%.⁸ The changes within industries are most important in increasing abstract and routine occupations' employment shares, while changes between industries matter the most for manual occupations.

⁸These percentages correspond to the relative contributions of Table 1.2 weighted by their employment changes.

The following paragraphs describe this shift-share decomposition. The basic idea comes from expressing the aggregate share of each occupation as a weighted average. For period t :

$$p_t(j) = \sum_I s_t(I)p_t(j|I) \quad (1.1)$$

where $p_t(j)$ is the economy-wide employment share of occupation j , $s_t(I)$ is the economy-wide employment share of industry I , and $p_t(j|I)$ is the share of occupation j in industry I .

The change between period 0 and t can be decomposed into a *between* industry effect, and a *within* industry effect. The between industry effect refers to structural transformation. Services is more intensive in abstract and manual occupations than goods, so higher employment shares in services imply higher abstract and manual occupational shares in the economy. The within effect refers to the occupational mix inside each industry. Over time, both the production of goods and services demand more abstract occupations, and less routine ones. In particular:

$$\Delta p_t(j) = \underbrace{\sum_I \Delta s_t(I)\bar{p}(j|I)}_{\text{Between industries effect}} + \underbrace{\sum_I \Delta p_t(j|I)\bar{s}(I)}_{\text{Within industries effect}} \quad (1.2)$$

$\bar{p}(j|I)$ is the average between time 0 and t of the conditional occupation share, and $\bar{s}(I)$ that of the industry share. Notice that this decomposition does not have any residual term, since the employment changes are evaluated at their average over the period.

Between-industry changes are the most important component for manual occupations, while within-industry changes are the stronger one for routine and abstract occupations. Table 1.2 shows this. About 90 percent of all of the increase in manual

Table 1.2: Shift-Share Decomposition of Changes in Occupational Shares

Occupation	Absolute Change (p.p.)			Relative Contribution (%)	
	Total	Between	Within	Between	Within
	Change	Industries	Industries	Industries	Industries
Manual	3.8	3.4	0.4	89.5	10.5
Routine	-16.9	-5.7	-11.1	33.7	66.3
Abstract	13.1	2.3	10.8	17.6	82.4

Source: author's calculations using CPS.

occupations is due to changes between industries, and over 80 percent of the increase in abstract occupations is due to changes within industries. Finally, changes within industries are the stronger component in the decrease of routine occupations. Overall, weighting these relative contributions by the magnitude of the employment share change reveals that 66% of polarization is due to changes within industries, and 34% due to changes between occupations.

In conclusion, the within-industry composition of the economy is the main driver of polarization. It is important to include the industrial composition, however, especially to account for the rise in manual occupations' employment share. This goes in line with previous findings: Autor and Dorn (2013) argue that the expansion of personal services lies behind the increase in low-skill, manual occupations, while Tüzemen and Willis (2013) and Bárány and Siegel (2018) also show that within-industry changes are more important for routine and abstract occupations. These also show similar results when using finer classifications in industries and occupations.

1.2.3 Lower Non-employment Filled the Edges

The “hollowing out” of the employment distribution can be easily associated to decreases in routine employment, pushing workers into lower paying jobs, and non-employment. In this section, I show that the data is at odds with that interpretation: if we include non-employment as an additional category, routine occupations remain fairly stable as a share of *total population*, as opposed to its decreasing employment share. In addition, the decrease in non-employment allowed for increases in manual and abstract occupations. This means that lower non-employment accounts for job polarization, not decreases in routine employment.

Notice that including non-employed as an additional category changes the reference point: now we study total population aged between 25 and 65 years, rather than employed workers. This allows to study the net flows among these categories, and to better understand the changes leading to polarization.

Figure 1.4 plots non-employment, and workers in manual, routine, and abstract occupations as shares of total population. There are three main takeaways from this figure: manual and abstract occupations increased their share over time, routine’s share remained fairly constant over time, and non-employment decreased for most of this period.

The share of manual and abstract occupations in total population increased during this period, the same way these increased their shares in employment. This means that the increase in the employment shares of manual and abstract occupations responds to a net inflow of workers into these occupations. During these years, the share of these occupations in total population grew by a factor of 1.46 and 1.69, respectively.

The share of routine occupations in employment mostly fell because of lower non-

employment: only 26% can be attributed to lower routine employment, and the remaining 74 is due to higher employment in manual and abstract occupations. For this, I use a Taylor approximation of the share in employment of routine occupations $p_t(r)$:


$$p_t(r) = \frac{R_t}{R_t + NR_t} \quad (1.3)$$

where R_t represents the share in total population of routine occupations, and NR_t represents the share in total population of non-routine occupations. Until 2007, only 9.7% of the decrease in routine occupation's employment share was due to a lower share of routine occupations in total population, and increased to 26% after the great recession of 2008.

Finally, these changes were possible because of the reduction in non-employment. Most of the decreases happened until 2000, and had an increase after the great recession that. However, the increases in employment in manual and abstract occupations, that are mostly behind polarization, are the main forces explaining job polarization.

To sum up, in this section I argue that lower non-employment accounts for job polarization, not decreases in routine employment. Net flows out of non-employment, therefore, are key to explain overall polarization. If job displacements were the main cause for the lower employment share in routine occupations, their share in total population would have decreased significantly over time. This is not the case, so we can discard displacements (or increases of it) as the main force behind the fall of routine occupations. This is consistent with the findings in Cortés (2016), where panel data show scant evidence of displacement.

Figure 1.4: Occupational Job Polarization & Labor Non-Participation



Shaded areas indicate years of major changes in occupational codes. These percentages refer to each category’s share in total population.

Source: author’s calculations using CPS.

1.3 A Labor Allocation Model

Section 1.2 presented a new stylized fact about job polarization: that most of it happened through decreases in non-employment. In this section, I use a model of labor allocation to study quantitatively how changes in non-employment affected it. Its contribution is to provide a general equilibrium framework to analyze the link between non-employment, structural transformation, and job polarization.

Broadly speaking, this is a model of structural transformation where consumers choose between goods and services; goods are produced in the market, and services can be either produced in the market, or at home. In addition, market production demands workers to perform different tasks, which translate into the three occupations studied in the last section. These margins are needed to account for the employment

shifts presented in the last section. The driving force is *task specific* technical progress, so the difference in their growth rates induces the three main results: polarization, structural transformation, and changes in non-employment.

The goal of this model is to assess the importance of changes in non-employment for job polarization. The link between these two comes from the consumption of services. Activities that were typically prepared within the household, which correspond to home services in the model, are going through a process of “marketization.” This means that they are now being traded in the marketplace, and some of the activities performed as home production are substituted for services produced in the market.⁹ This marketization impacts the demand of occupations: higher demand for market services boosts the demand for manual and abstract occupations, while higher demand for home services increases non-employment. This trade-off between market and non-market activities has not gone unnoticed in the literature, but few have focused it towards occupations. The following discussion presents only the relevant aspects of the model; the complete presentation is left to Appendix A.3.

1.3.1 Preferences & Technology

This is a discrete-time model where time runs forever. On the consumption side, there are identical households of measure one. These households value three types of consumption: goods that can only be bought in the market, services bought in the market, and services produced at home, as in Ngai and Petrongolo (2014). Differently from them, households only value consumption. The utility level in period t is

⁹An early recognition of the importance of home production, and the tendency towards marketization is present in Kuznets (1941).

aggregated according to a nested CES specification:

$$U_t(C_{Gt}, C_{St}) = \left[\omega_G^{\frac{1}{\varepsilon}} (C_{Gt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_S^{\frac{1}{\varepsilon}} (C_{St})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \quad (1.4)$$

where C_{Gt} and C_{St} denote the consumption of goods, and a basket of services. Their relative preference weights are ω_G and ω_S , which add up to one, and individually are between zero and one. These two consumption categories consist of broad, and disparate types of consumption, which have an elasticity of substitution of $\varepsilon > 0$.

The basket of services is also represented through a CES aggregator:

$$C_{St} = \left[\varphi_M^{\frac{1}{\eta}} (C_{MSt})^{\frac{\eta-1}{\eta}} + \varphi_H^{\frac{1}{\eta}} (C_{HSt})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} \quad (1.5)$$

where C_{MSt} and C_{HSt} denote market and home services. Their relative weights are φ_M and φ_H , which also add up to one, and individually are between zero and one. The types of services provided by these two are much more similar, and their elasticity of substitution is $\eta > 0$.

The production side of this economy has three types of output: goods, market services, and home services. Firms producing goods and market services combine manual, routine, and abstract tasks, while households producing their own services require home workers. Notationally, I denote the sector of production by uppercase letters, so $I \in \{G, MS, HS\}$. Occupations and tasks are denoted by lowercase letters, so $j \in \{m, r, a, h\}$ stand for manual, routine, abstract, and home.

Firms producing goods and market services follow the task approach to production, as in Acemoglu and Autor (2011). Firms, then, combine *tasks* to obtain output, and hire workers to produce those tasks. Ultimately, these are combined through a CES production function:

$$Y_{It} = \left[\sum_{j \in \{m, r, a\}} \alpha_{Ij}^{\frac{1}{\sigma}} (A_{jt} N_{Ijt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (1.6)$$

where N_{Ijt} denotes the labor input that industry I uses to produce task j , and A_{jt} is the labor productivity in task j . $A_{jt}N_{Ijt}$ then denotes the task input of j in industry I . This task approach follows the same grouping principle for the occupations in the empirical section, therefore workers hired to perform task j are working in occupation j . The elasticity of substitution between the task inputs is $\sigma > 0$, and $\alpha_{Ij} \in (0, 1)$ is the intensity of task j in industry I . Notice that productivity is task-specific, and is the same across industries, as in Duernecker and Herrendorf (2016). Differently to them, this technology is defined over three tasks.

Home services only require home workers to produce, so its technology is represented by a linear function:

$$Y_{HSt} = A_{ht}N_{HSt} \tag{1.7}$$

where A_{ht} denotes the productivity of home workers, and N_{HSt} denotes the labor input used in home services. These home workers provide the model's counterpart to non-employment.

Labor is homogeneous and perfectly mobile. This means that any two workers can switch between occupations costlessly, and be equally productive if working on the same task. In addition, I normalize the mass of workers to 1, so that all labor inputs in the model match their empirical counterparts as shares in total population.

1.3.2 *Competitive Equilibrium Outcomes*

In this section, I analyze the labor reallocation patterns through the lens of the model, assuming perfect competition. Over time, these will only change due to the task-specific productivities. Therefore, I will analyze the productivity changes that must take place to reproduce the employment patterns in Section 1.2. The model structure requires productivity growth to be the highest in routine tasks, followed by

manual, abstract and home production.

This result hinges on three key assumptions related to preferences and technology. In particular, I assume:

1. Tasks to be complements in production.
2. Home and market services to be substitutes in consumption.
3. Goods and services to be complements in consumption.

These are all standard assumptions in the literature, and I will discuss these choices in the following subsections. Intuitively, complementarity means that the tasks and categories are very different, and are hard to substitute from one another. The opposite is the case for substitutes: similar categories that are easier to exchange from one another.

Under competitive markets, firms will maximize their profits, and households their utility functions as price takers, as usual. I will analyze the result of these trade-offs around four components: job polarization, industrial productivity, structural transformation, and changes in non-employment.

Explaining Job Polarization

Job polarization within industries results from firms' optimal combination of inputs. Under perfect labor mobility, wages equalize across occupations. Therefore, the demand for workers depends on the productivities of each of the tasks, and their elasticity of substitution.

This is expressed more clearly using firms' first order conditions. In industry I , equalizing the marginal revenue from workers in tasks j and k yields the relative

demand of occupations:

$$\frac{N_{Ijt}}{N_{Ikt}} = \frac{\alpha_{Ij}}{\alpha_{Ik}} \left(\frac{A_{kt}}{A_{jt}} \right)^{1-\sigma} \quad (1.8)$$

This relative occupational demand depends on the relative intensities in that particular industry, and the relative task productivities. The elasticity of substitution, σ , dictates how strongly industries react to changes in productivities.

When occupations are complements, rather than substitutes,¹⁰ firms react to productivity increases by demanding more of the other occupations. Equation (1.8) shows this more clearly: complementarity implies $\sigma < 1$, and increases in the productivity of task j will boost the relative demand of occupation k . This is due to the cross-productivity effects. If task j increases its productivity, the marginal productivity of the workers in other occupations increases by more, and creates an “excess supply” in occupation j . To benefit the most from this productivity boost, firms demand more of the other occupations.

Recall from Figure 1.2 that polarization took place within industries. To reproduce that pattern, the growth rate in the productivity of routine tasks has to be the highest one, followed by manual, and finally abstract. This is due to the complementarity assumption; it would be reverted if occupations were substitutes.

The productivity terms are modeled in reduced-form way, so these encompass many forces pushing down the demand of routine occupations. Factors affecting these productivities include plain capital accumulation, capital deepening as in Acemoglu and Guerrieri (2008), and increased substitutability with routine workers. The

¹⁰Recall that in this CES formulation, a unitary elasticity of substitution ($\sigma = 1$) implies a Cobb-Douglas production function. As their substitutability decreases, and $\sigma < 1$, inputs are called gross complements. In the limit case, when $\sigma \rightarrow 0$, the function converges to a Leontief technology. When $\sigma > 1$, inputs are gross substitutes, given their higher degree of substitutability. When $\sigma \rightarrow \infty$, substitutability is perfect and the production function converges to a linear technology.

common denominator in these changes is diminishing the cost of performing routine tasks over time. Goos *et al.* (2014) use a similar approach, but focusing on the occupational cost functions. Their approach on input prices is observationally equivalent to one where the production function is modeled explicitly, as I do in this model.

Treating tasks as complements amounts to assuming that their elasticity of substitution, σ , is less than one. Most of the studies using different labor inputs treat these as substitutes, rather than complements.¹¹ The task approach to occupations calls for a different interpretation. Since these tasks are very different, it is difficult to substitute for one another. One could think of a production process that combines manual, routine, and abstract tasks, and does so in relatively fixed proportions. This technology is closer to a Leontief specification, which can be represented by a CES production function with a low elasticity of substitution. The effect of increases in productivities, as in Acemoglu and Autor (2011), is to create an “excess supply” of that task, holding the labor composition constant.

Lastly, since productivity changes happen at the task level, it affects all industries equally. This is also observed in the data. In the model, when the elasticity of substitution is equal over industries, the changes in occupational demands should be the same. In the data, comparing the changes in relative occupational shares reveals a similar relationship.¹²

¹¹Examples of this are Katz and Murphy (1992), that distinguish between educated and non-educated workers, and Caselli (2015), that distinguishes between experienced and inexperienced workers.

¹²One way to test this relationship is to take the ratio of (1.8) for a given occupation pair over industries. This shows little variance over time: the coefficients of variation are 0.21 and 0.11, which indicate relative stability.

Explaining Industrial Productivity

The second pattern to focus on is the evolution of labor productivity at the industry level. For the production of goods and market services, occupational demands can be aggregated to a linear technology in total industry demand. At the optimal occupational demands, the production function (A.2) can be rewritten as:

$$Y_{It} = \tilde{A}_{It} N_{It} \quad (1.9)$$

where

$$\tilde{A}_{It} = \left[\alpha_{Im} \left(\frac{1}{A_{mt}} \right)^{1-\sigma} + \alpha_{Ir} \left(\frac{1}{A_{rt}} \right)^{1-\sigma} + \alpha_{Ia} \left(\frac{1}{A_{at}} \right)^{1-\sigma} \right]^{\frac{-1}{1-\sigma}} \quad (1.10)$$

is the average labor productivity in industry $I \in \{G, M\}$, and N_{It} is its associated labor demand.

This productivity is a weighted average of each task's productivity, when labor inputs are combined optimally. With different productivity growth rates in tasks, growth at the industry level will be non-linear, will vary over industries, and will depend on each industry's occupational intensity.

When tasks are complements, and with constant productivity growth, the industry that uses more intensively the task with the highest (lowest) productivity growth increases its overall productivity the most (least). Asymptotically, industry productivity growth rates will converge to the rate of the task with the lowest growth rate. As the share of the occupation with the higher productivity growth decreases, so does its contribution to the growth rate of that industry's productivity. Polarization, then, ends up dampening these productivity gains. At some point in the reallocation process, the share of the occupation with the lowest productivity growth will be so high that its effect on industry productivity will be the only discernible one.

Baumol's cost disease lies at the heart of these dynamics. In their reappraisal of the unbalanced growth model, Baumol *et al.* (1985) discuss that the progressivity or

stagnancy of economic activities is caused by the technological advances behind their inputs, which correspond to occupations in this setting. The literature in structural transformation has established, by several measures, that goods-producing industries have had higher productivity growth, compared to services-producing industries.¹³ This would require goods to be more intensive in routine tasks, and market services to be more intensive in abstract tasks. The data show that this is clearly the case: on average, goods use 63% more routine workers than services, and services use 50% more abstract.

For notational completeness, define the industry equivalent of (1.10) for home services production:

$$Y_{Ht} = \tilde{A}_{Ht} N_{Ht} \quad \text{where} \quad \tilde{A}_{Ht} = A_{ht} \quad (1.11)$$

Marketization

The third pattern to focus on is the marketization of home production, which relates to the reallocation of productive resources from the home sector to the market.

The decision to consume home services is slightly different than that of goods and market services. To consume home services, households must produce them themselves and give up the market income that they would otherwise earn. The opportunity cost of home production is $p_{Ht} = w_t / \tilde{A}_{Ht}$. In equilibrium, the relative price of home to market services will be inversely related to the sectoral productivities:

$$\frac{p_{Ht}}{p_{Mt}} = \frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}} \quad (1.12)$$

Households decide their consumption patterns in services taking into account these

¹³A review of this literature, and evidence for several countries is presented in Herrendorf *et al.* (2014a).

relative prices. This translates into the following labor allocations:

$$\frac{N_{Ht}}{N_{Mt}} = \frac{\varphi_H}{\varphi_M} \left(\frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}} \right)^{1-\eta} \quad (1.13)$$

Then, again, the evolution over time of this labor allocation will depend on the degree of complementarity and the relative productivity growth. If home and market services are good substitutes, increases in the relative price of home production lead households to substitute its consumption with market services. To do this, they decrease the relative amount of labor dedicated to home production.

This is the marketization result discussed in Freeman and Schettkat (2005): the United States has seen a shift of traditional household production to the market. In the model this would reflect a decrease in the relative labor allocated to home production. This, again, calls upon questioning how reasonable the assumptions behind this result are.

Firstly, we should analyze the existence of the home sector itself. Kuznets (1941) was well aware of this fact, and pointed out that incomes within the family economy were a prominent missing item in his estimates of national income. His approximations amounted to more than a quarter of national income in 1929. Hill (1985) uses time-use surveys to establish that for married couples in the mid-1970s, time spent on home work was only slightly behind market work. More recently, Aguiar and Hurst (2007) also find that the amount of hours in home production are substantial, compared to market hours, albeit decreasing over time. These observations should be enough to agree with Benhabib *et al.* (1991, p. 1185): “models without home production implicitly make the assumption that the willingness or the incentive of individuals to substitute between market and nonmarket activity is small, but this does not seem to be the conclusion one would want to draw from the evidence.”

Secondly, in this model, the time that is not spent in market work is dedicated

exclusively to home production. Ngai and Pissarides (2008b) posit that in terms of the production carried out in the household, only home services remain. They analyze a much longer time period, and use a structure where home work could be devoted to the three sectors in their study: agriculture, manufacturing, and services. By the late 1920s, they conclude that home production in agriculture and manufacturing was practically gone. Thus, I align with their observations, and assume these productions away.

In this model, the people that are not working in the market are engaging in home production. Time use surveys show that non-employed people engage mostly in home production, performing activities like housework, cooking, and child care (Ngai and Pissarides, 2008b). Many of these activities have close counterparts in the market, particularly in the services sector. Other articles building on this idea are Buera and Kaboski (2012) and Ngai and Petrongolo (2014), which I adopt as well.

Thirdly, assuming market and home services are good substitutes in consumption ($\eta > 1$) should not come as a controversial issue. Housework, shopping, food preparation, and caring for other people are among the activities that take most of the time in home production, according to time use surveys. These are all activities that can be easily purchased in the modern marketplace, thus their high degree of substitutability.

Lastly, since market and home services are good substitutes, households will tilt their consumption to the sector with lower price, which is the sector with the higher productivity growth rate. The increase in market participation requires a considerable difference between the growth of home and market productivities. Bridgman (2016) presents evidence to suggest that, effectively, productivity growth in the market has outpaced home productivity, in particular during this paper's period of study.

Structural Transformation

The fourth pattern focuses on the reallocation of consumption and productive resources between market industries, in particular from goods to services. Their relative prices are, again, inversely related to their sectoral productivities:

$$\frac{p_{Gt}}{p_{Mt}} = \frac{\tilde{A}_{Mt}}{\tilde{A}_{Gt}} \quad (1.14)$$

Preferences are homothetic, so there are no income effects. The expenditure ratio and labor allocation between market goods and services consists of two parts: one that is a price effect, and another one that is a marketization effect. These are represented by:

$$\begin{aligned} \frac{p_{Mt}C_{Mt}}{p_{Gt}C_{Gt}} &= \frac{N_{Mt}}{N_{Gt}} \\ &= \underbrace{\frac{\omega_S}{\omega_G} \left(\frac{\tilde{A}_{Gt}}{\tilde{A}_{Mt}} \right)^{1-\varepsilon}}_{\text{Price effect}} \underbrace{\left\{ \varphi_M^{1-\varepsilon} \left[1 + \frac{\varphi_H}{\varphi_M} \left(\frac{\tilde{A}_{Mt}}{\tilde{A}_{Ht}} \right)^{1-\eta} \right]^{\eta-\varepsilon} \right\}^{\frac{1}{1-\eta}}}_{\text{Marketization effect}} \end{aligned} \quad (1.15)$$

The price effect responds to relative productivity between *market* industries, while the marketization effect responds to relative productivity between *service* industries.

The price effect behaves similarly to the canonical structural transformation model of Ngai and Pissarides (2007b). In this setting, however, task-specific productivity growth is responsible for the growth at the industry level. Complementarity between goods and services (which requires $\varepsilon < 1$) implies that increases in the relative price of services result in increases in its expenditure share.

The marketization effect, on the other hand, is somewhat similar to an income effect. An income effect would induce different consumption patterns, holding constant the productivities across occupations. Within the context of structural transformation, Kongsamut *et al.* (2001a) introduce income effects with Stone-Geary preferences, and interpret the non-homotheticity term as home production. In this model,

as home services become comparatively more expensive, households make up for this by switching out of home production, increasing market work and purchasing more services in the market.

Non-employment

The last component to focus on is the net effect of these forces on non-employment. Structural transformation, with its decrease in the relative price of goods, makes households want to increase their consumption of services. Marketization, on the other hand, makes home production relatively more expensive to market services. Then, the non-employment decision involves a trade-off between home production and market consumption. In equilibrium, the ratio of market employment to non-participation is:

$$\frac{N_{Gt} + N_{Mt}}{N_{Ht}} = \underbrace{\left[1 + \frac{\omega_G}{\omega_S} \left(\frac{\tilde{A}_{St}}{\tilde{A}_{Gt}} \right)^{1-\varepsilon} \right]}_{\text{Structural transformation}} \underbrace{\left[1 + \frac{\varphi_M}{\varphi_H} \left(\frac{\tilde{A}_{Ht}}{\tilde{A}_{Mt}} \right)^{1-\eta} \right]}_{\text{Marketization}} - 1 \quad (1.16)$$

This expression separates the forces of structural transformation and of marketization. An increase in this ratio implies an increase in labor force participation.

Structural transformation frees up labor that can be used to produce services. This, absent a strong reallocation of consumption within services, would imply higher non-employment. Home production would end up filling part of this increased demand for services. Marketization has the opposing effect, since it becomes relatively cheaper to consume more services from the market.

In the data, there is a sizable decrease in non-employment, most of which is used to fill the increased demand for abstract occupations. This means that the forces of marketization are considerably stronger than those of structural transformation.

Summary of the Model

A brief summary of this model starts with productivity growth rates at the task level, since their differences are the source of the reallocation patterns. The growth rate is higher in routine tasks, followed by manual, abstract, and finally home production. The interaction between these productivities, technology, and preferences yields five outcomes:

Polarization: firms demand more workers in the abstract and manual occupations, relative to routine, because tasks are complements in production.

Industry Productivity Growth: productivity growth is higher in goods because it is more intensive in routine tasks than market services.

Marketization: households work more in the market and substitute home with market services because of increasing opportunity costs of home production.

Structural Transformation: households demand more services because the relative price of goods decreases.

Non-employment: non-employment decreases because the effect of marketization dominates the effect of structural transformation on services.

1.4 Quantitative Results: How Non-employment Affects Polarization

This section explores the quantitative side of the model. First, I describe the estimation procedure, and analyze its results. With the estimation in hand, I conduct two counterfactual exercises to assess the importance of the increase in labor force participation.

1.4.1 Calibrating the Model

In this section, I explain briefly how to calibrate the model, and the moments I use. To begin with, I choose the two elasticities in the utility function based on previous studies. Based on Duernecker and Herrendorf (2016), I set $\varepsilon = 0.05$ (between goods and compound services). Based on Rogerson (2009) and Ngai and Petrongolo (2014), I set $\eta = 2.3$ (between market and home services).¹⁴

With these restrictions, there are eleven time-invariant parameters and four terminal conditions to determine: six task intensities (three for each market industry), the labor elasticity of substitution in market production, the four final task productivities, and the four preference weights in the consumption (for goods and services, and for home and market services). Assuming constant growth rates, there is only need to look at the initial and final years, which are denoted here by $t = 0$ and $t = T$.¹⁵ I back out their estimates from US data following these steps:

1. Impose the normalization $A_{m0} = A_{r0} = A_{a0} = A_{h0} = 1$.
2. Use the initial market occupation shares N_{Ij0} to solve for α_{Ij} .
3. Use the initial home and market services shares N_{H0}, N_{M0} to solve for φ_H and φ_M .
4. Use the initial market goods and services shares N_{G0}, N_{S0} to solve for ω_G and ω_S .

¹⁴This elasticity is in the high end of the estimates available. It was purposely chosen as such, because these come from studies looking at the substitution between home and *total* market goods. This selection attempts to make up for goods being included.

¹⁵Notice this same procedure could be used to infer a more detailed productivity path on a yearly basis, but the smooth labor share paths suggest this is a reasonable assumption.

5. Use the final employment shares in the market services industry N_{MjT} to solve for final relative productivities $(A_{aT}/A_{rT})^{1-\sigma}$ and $(A_{mT}/A_{rT})^{1-\sigma}$.
6. Use the final relative employment share N_{MT}/N_{ST} to solve for the labor elasticity of substitution σ .
7. Use the growth factor of real per capita GDP to solve for A_{rT} .
8. Use the final home and market services shares to solve for A_{hT} .

Further details of this procedure are discussed in Appendix A.4. Table 1.3 shows the time-invariant parameters of the model, and Table 1.4 the occupation and industry productivity estimates.

These results are in line with those explained in the discussion section, so the qualitative predictions remain. Now we can comment on their quantitative side. As expected, growth in all occupation-specific productivities is positive. Productivity growth is such that by 2018, a worker in routine occupations is 35% more productive than a worker in manual occupations, and almost twice as much than a worker in abstract occupations. This goes in line with the routinization hypothesis, but established as a force working since (at least) the beginning of the sample. The elasticity of substitution between occupations is considerably lower than other estimates. Again, this stems from considering occupations as different factors of production, which is induced by the task-oriented grouping.

Table 1.5 shows the model's predictions with the estimated parameters. By design, it is able to match all of the labor shares in 1968, and the relative labor shares within services (for the three occupations in market services, plus the ratio between market and home services). The model is fairly successful at reproducing the relative occupation shares in the goods industry, and a little less so in reproducing the decrease

Table 1.3: Model Parameters

	Parameter	Value	Source
ε :	Elasticity of substitution in utility b/w goods and combined services	0.05	Herrendorf <i>et al.</i> (2013)
η :	Elasticity of substitution in utility b/w home and market services	2.30	Ngai and Petrongolo (2014)
ω_G :	Preference weight for goods in utility	0.27	Initial goods and services relative price
ω_S :	Preference weight for combined services in utility	0.73	Initial goods and services relative price
φ_H :	Preference weight for home services in utility	0.45	Initial home and market services labor shares
φ_M :	Preference weight for market services in utility	0.55	Initial home and market services labor shares
σ :	Labor elasticity of substitution in production b/w market occupations	0.22	Final employment shares in market services
α_{Ga} :	Intensity of abstract occupations in market goods production	0.22	Initial industry-specific occupation shares
α_{Gr} :	Intensity of routine occupations in market goods production	0.77	Initial industry-specific occupation shares
α_{Gm} :	Intensity of manual occupations in market goods production	0.01	Initial industry-specific occupation shares
α_{Ma} :	Intensity of abstract occupations in market services production	0.35	Initial industry-specific occupation shares
α_{Mr} :	Intensity of routine occupations in market services production	0.47	Initial industry-specific occupation shares
α_{Mm} :	Intensity of manual occupations in market services production	0.18	Initial industry-specific occupation shares

All parameters, except the first two elasticities, are estimated from CPS data. See Section 1.4.1 and Appendix A.4 for more details.

Table 1.4: Productivity Estimates

	Year		Average
	1968	2018	Growth Rate
A_h : Home production	1	1.01	0.03%
A_a : Abstract occupations	1	1.32	0.56%
A_r : Routine occupations	1	2.59	1.96%
A_m : Manual occupations	1	1.93	1.35%
\tilde{A}_G : Market goods industry	1	2.15	1.57%
\tilde{A}_M : Market services industry	1	1.87	1.29%

in the goods share. The normalized root-mean-squared-error (RMSE) shows that for a model with constant growth rates, it is still satisfactory in capturing the main forces driving the changes in labor markets.

1.4.2 Assessing the Importance of Non-employment

One of the main goals of this paper is to assess the importance of lower non-employment in the distribution of employment in the United States. The model at hand allows us to do this, and study its effects on structural transformation, on overall polarization, and on sectoral output. To do that, I study two experiments. In the first one, I restrict the model so that non-employment stays at its 1968 level. In the second one, I increase the productivity growth in home production to shut down the marketization channel. The comparison point of these exercises are the output levels and labor allocations of the baseline model in 2018. The first year of my sample coincides with the years when labor force participation started to increase. Juhn and Potter (2006) report that between 1948 and 1968 the participation rate remained

Table 1.5: Actual and Fitted Labor Shares

	1968		2018		Normalized
	Data	Model	Data	Model	RMSE
Share of total population in:					
Goods production	27.3	27.3	16.3	20.9	15.0%
Market services production	40.1	40.1	60.7	58.7	7.5%
Home services production	32.6	32.6	23.0	20.4	8.7%
Share of goods labor demand in:					
Manual occupations	1.2	1.2	1.0	1.3	18.9%
Routine occupations	76.7	76.7	70.5	65.6	1.9%
Abstract occupations	22.1	22.1	28.5	33.2	8.1%
Share of market services labor demand in:					
Manual occupations	18.1	18.1	17.7	17.6	8.7%
Routine occupations	47.4	47.4	38.3	36.2	8.0%
Abstract occupations	34.5	34.5	43.9	46.2	5.3%

Normalized root-mean-squared-error is calculated using the average of the time series.

relatively stable, and after that it increased. This provides a convenient turning point to perform counterfactual exercises.

Freezing Non-Employment

The expansion of the labor force, which has lowered non-employment, has drawn a considerable amount of attention recently, and has been associated with the insertion of women into the workplace. Several interpretations have been presented to explain this. One of them focuses on the social attitudes towards women's work: Fernández

et al. (2004) discuss the gradual transformation of the family model (where working mothers set an example for future generations), while Goldin (2006) discusses several other social changes fueling the “quiet revolution” that made women think in terms of lifelong careers. In the spirit of these articles, I freeze non-employment at its level in 1968. One interpretation of this counterfactual would be to consider how different the productive structure would have been, had societal attitudes towards women’s work not changed as much. The results of this are in Table 1.6.

Table 1.6: Freezing Non-employment, Counterfactual Results

	Baseline	Counterfactual	Change
	model	prediction	(%)
Non-employment	0.20	0.33	59.4
Goods output	0.46	0.45	-1.8
Market services output	1.08	0.87	-19.2
Share in labor force of:			
Goods industry	0.26	0.30	15.5
Services industry	0.74	0.70	-5.5
Share in total population of:			
Manual occupations	0.13	0.13	-5.0
Routine occupations	0.44	0.45	2.7
Abstract occupations	0.43	0.42	-1.2

This are the model’s predictions for 2018. The counterfactual exercise freezes the labor force participation rate at its 1968 level, without assuming different rates of productivity growth. Percent changes are reported with respect to the baseline model’s predictions.

Holding non-employment fixed decreases output in both market sectors, disproportionately so for market services. This is not surprising, since this experiment deliberately holds down the inputs for market production, and restricts households to consume more home services than they would otherwise want. Within the workforce, this slows down structural transformation: the employment share of goods ends up being 16% higher, despite market productivities remaining the same. This division of labor between goods and services makes the productive structure look like the model's prediction for 1999. Focusing on job polarization, the forces acting *within* both industries would still take place, so the labor force would polarize to a similar degree. The main difference is that in this case, the entirety of the adjustment takes place *within* the labor force, instead of through an expansion of it. Then, worker displacement would be a significant source of adjustment.

Increasing Home Productivity

Another explanation for the change in non-employment focuses on the productivity of the household sector. In this vein, the second experiment keeps non-employment from decreasing, but because of very different reasons. Bridgman (2016) suggests that productivity growth in the home sector had a considerable slowdown, coinciding with the initial years of this CPS sample. He also reports that previous to these years, home productivity had been growing at a similar pace to market activities. This is consistent with non-employment being stable before 1968, and increasing after this break. This begs the question: if the slowdown in home productivity was the sole cause of marketization, what would have happened without it? That is, how different would the productive structure look like if productivity in the home production sector had grown at the same rate as market services? The outcome of this experiment is in Table 1.7.

Table 1.7: Increasing Home Productivity, Counterfactual Results

	Baseline	Counterfactual	Change
	model	prediction	(%)
Non-employment	0.20	0.34	65.3
Goods output	0.47	0.56	18.2
Market services output	1.14	0.81	-29.2
Share in labor force of:			
Goods industry	0.26	0.37	42.0
Services industry	0.74	0.63	-14.9
Share in total population of:			
Manual occupations	0.13	0.12	-13.5
Routine occupations	0.44	0.47	-7.4
Abstract occupations	0.43	0.41	-3.4

This are the model's predictions for 2018. The counterfactual exercise shuts down the marketization channel by increasing home production's productivity at the same pace as market services. Percent changes are reported with respect to the baseline model's predictions.

The results in terms of output are more dramatic in this case: production in goods increases, while in market services it decreases even more. Households are much more efficient in producing their own services, so they do not increase their reliance on the market. This explains the greater hit that market services take. Non-employment increases; since I am purposely shutting down the marketization channel, only structural transformation has an effect on the consumption of total services (equation (1.16) shows this). These effects add up to goods having an even

higher share in market employment, which increases by almost 40%. The resulting division of labor makes the productive structure look like the model's prediction for 1977. In terms of polarization, the intra-industry reallocation would still take place. In this case, however, lower growth in market services leads to less polarization than in the first counterfactual. As in the first counterfactual, all the reallocation would happen within the labor force.

These two exercises illustrate how important the decrease in non-employment is for both structural transformation and job polarization. In the two cases, both structural transformation and polarization are slowed down by the induced dynamics of home production. This raises the question of, for instance, to what extent the differences in the labor market structure of the United States and some European countries are driven by the channels highlighted in this model. Prescott (2004) explores the issue of differences in hours worked by looking at taxes. Home production could play a significant role in this setting. If cultural preferences favor home production, despite a lower productivity in it, it could look more like the first counterfactual exercise. If, on the other hand, European households did not experience a significant decrease in their home productivity growth, it could look more like the second counterfactual. These two cases show very different implications for growth and the causes for the differences in their productive structures. Polarization slows down, but the road it takes is quite different. Unfortunately, this model is unable to speak of mobility costs by design. Displacement costs would play an important role in terms of welfare, and could possibly slow down the adjustment process. This interesting avenue of research is left for other research projects.

1.5 Conclusions

In this article I study three recent trends in the United States labor market: job polarization, structural transformation, and decreases in non-employment. With the goal of quantifying the importance of the non-participation margin, I propose a labor allocation model to explain this occupational and industrial structure.

Quantitatively, the model implies higher growth in the occupation of routine occupations, followed by manual, abstract, and finally home production. It is able to reproduce the occupational structure within industries, and the shift towards market services.

Counterfactual exercises suggest that this expansion is very important. Holding constant the attitudes that allowed lower non-employment decreases the production of goods and market services by 2 and 19%, holds back structural transformation to its 1999 level, and decreases polarization by an average of 3%. The second exercise, inspired by the home productivity slowdown reported in Bridgman (2016), has home productivity growing at the rate of market services. This increases the production of goods by 18%, decreases the production of market services by 29%, holds back structural transformation to its 1977 level, and decreases polarization by an average of 8%.

Future work includes extensions to this model to incorporate the age structure of the economy through an overlapping generations setting, considering the gender component to explain the wide differences in their labor market outcomes, and inducing labor heterogeneity to include an analysis of wage polarization.

Chapter 2

STRUCTURAL TRANSFORMATION BY COHORT

2.1 Introduction

One of the key stylized facts of economic growth is that it involves structural transformation: the reallocation of economic activity in predictable ways among the broad industries of an economy. Whereas poor countries typically produce and consume a high share of agriculture, growth entails a shift towards first manufacturing and then services.¹ A recent literature has explored different forms of preferences and technological progress that can generate this predictable reallocation of economic activity as a consequence of growth (Kongsamut *et al.*, 2001b; Ngai and Pissarides, 2007a).

Although the existing literature has advanced our understanding of structural transformation along many dimensions, it is largely silent about the interaction between structural transformation and labor markets, for two reasons. First, there are few stylized facts about structural transformation and labor markets.² Most papers use aggregate data from national accounts, which does little to clarify which workers are responsible for labor reallocation. Second, most papers focus on the special case of frictionless labor markets, which often allows for elegant analytical solutions but abstracts from the interactions we are interested in. Our goal is to make progress

¹See for example Schultz (1953) and Echevarria (1997) for early references, or Herrendorf *et al.* (2014b) for a recent overview. Herrendorf *et al.* (2014b) shows that structural transformation is a predictable function of PPP GDP p.c.

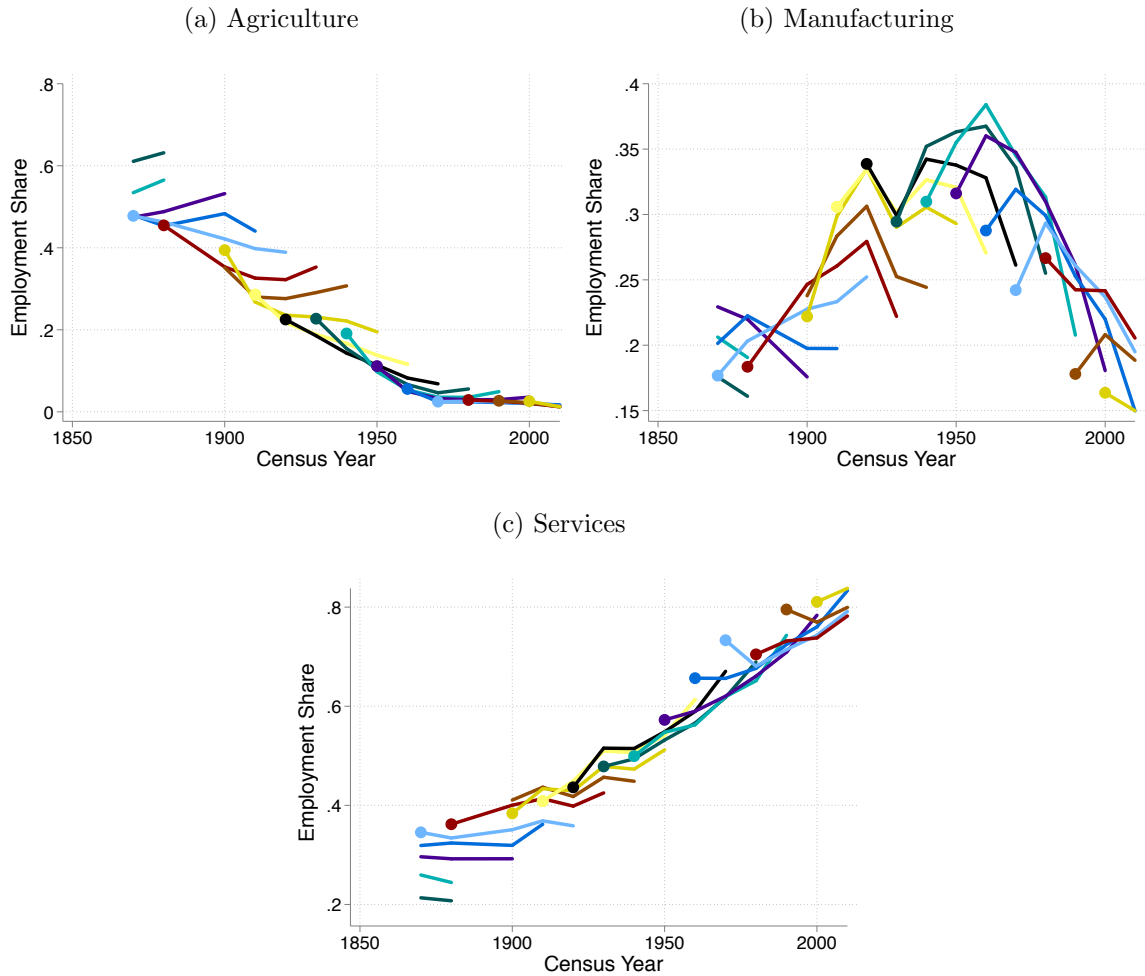
²There is however a large and related literature on gross worker flows between industries. See for example Kambourov and Manovskii (2008) for the U.S. and Carrillo-Tudela *et al.* (2016) for the United Kingdom.

on both fronts: we document new stylized facts of what sorts of workers reallocate during structural transformation; we develop a model consistent with these findings and use it to help understand the relationship between labor markets and structural transformation.

The starting point for our empirical contribution is to document stylized facts about which workers reallocate across sectors during structural transformation. To do so we utilize nationally representative, repeated cross sections spanning a long time series for the United States and shorter time series for 60 other countries. By using repeated cross sections we can track reallocation based on observable characteristics such as education. Figure 2.1 gives a visual representation of our approach for the case of the United States, 1870–2010. The three panels plot the employment shares by birth cohort for each decadal census in which these workers were between the ages of 20 and 70. Each line plots the time series for the workers in a particular birth cohort with the dot at the beginning of the line showing their employment share when they enter the workforce.

The overall pattern in the figure clearly shows the decline in the employment share of agriculture, the rise and then fall of the manufacturing share, and the increasing share of services over time. The second important point of this figure is that the lines for the individual cohorts do not overlap. Within a given year, newer and older cohorts have different employment patterns. In particular, the lines for each cohort appear to be “flatter” than the pattern for the overall economy. This is a very informal way of saying that within-cohort shifts in employment shares tend to be smaller than those in the overall economy and that differences in employment shares between cohorts are an important part of the sectoral reallocation of labor that has occurred in the United States. We formalize this finding using an accounting decomposition, which shows that 53 percent of reallocation happens between cohorts, both in the United

Figure 2.1: Structural Transformation in the United States



States and in our international sample.³ We also show that much of the within cohort (life-cycle) reallocation happens at earlier ages.

These findings suggest that new cohorts play a central role in the process of structural transformation. They lead us to formulate a heterogeneous agent overlapping

³Earlier authors have documented similar patterns for specific cases: Kim and Topel (1995) in Korea and Perez (2016) for Argentina. Porzio and Santangelo (2017) document similar facts for a large set of countries similar to ours for reallocation out of agriculture. We adopt the format of figure 2.1 from their work.

generations model of life-cycle career choice under switching costs and integrate it into a canonical model of structural transformation.⁴ The idea is that switching costs prevent older workers from moving between industries and hence give a prominent role to new cohorts for generating structural transformation. In doing so, we deviate from the common assumption of a single frictionless labor market that is made in many growth models, including those with structural transformation. In such a labor market, wages are equated across sectors, which runs counter to empirical evidence (Kim and Topel, 1995; Herrendorf and Schoellman, 2018).

An important contribution of our modeling strategy is to show how to formulate this problem in a tractable way. This is challenging at three levels. First, we need a tractable way to characterize the life-cycle career path of each individual agent as a function of wages. Second, wages themselves depend on the labor supply of past and future cohorts, which requires us to find a way to iterate on the entire path of labor supply and wages. Finally, our labor markets are part of structural transformation, which implies that the economy is experiencing unbalanced growth. We show how to overcome these challenges in our quantitative implementation.

At the aggregate we formulate structural transformation as in Ngai and Pissarides (2007a). Differential technology growth across sectors and a low elasticity of substitution in the utility function generate trends in both relative prices of goods produced in different sectors as well as the relative levels of labor demand across sectors. This formulation of structural transformation is useful for our purposes because it relies on homothetic preferences, which implies that we can solve for relative consumption as a function of only relative prices and not the entire distribution of income.

⁴We use the life-cycle career choice to mean the sequence of industries of employment. Duernecker and Herrendorf (2017) document a close linkage between reallocation across industries and reallocation across occupations during structural transformation.

Our main theoretical contribution is to integrate the life-cycle career choice of workers. Workers decide on their labor supply taking into account their idiosyncratic sector-specific skills, as in Lagakos and Waugh (2013) and in Barany and Siegel (2018), the current and future wages in each of the sectors, and the current and expected future retraining costs, as in Caselli and Coleman (2001b), associated with changing sectors of employment.⁵ We deviate from previous authors by formulating the career choice problem as a dynamic discrete choice problem. Doing so allows us to utilize known closed form solutions for life cycle sector choice and labor supply and avoid numerical integration, which greatly reduces the scale of the problem. Finally, we show how to adopt the extended path method of Fair and Taylor (1983) to this environment with unbalanced growth and solve for the equilibrium path of our model.

Our model provides four important insights. The first is that retraining costs for workers *accelerate* structural transformation. The reason for this counterintuitive result is the following. Given an initial sectoral allocation of labor, retraining costs slow sectoral reallocation down. However, forward-looking workers who face training and retraining costs change their initial labor allocation and shift their labor supply towards growing industries in anticipation of the future productivity and wage growth. This shift in initial sectoral choice more than compensates for lower life cycle sectoral reallocation.

The second result is that retraining costs need to increase with age to match cohort career profiles. To understand this, note that the model embeds an option value to working in a growing industry because of expected future relative wage growth. However, this option value declines as workers get older and have a lower

⁵Cociuba and MacGee (2018) also consider sectoral adjustment costs of workers, but do so in a stationary model with search frictions that is suitable at business cycle frequencies but does not allow for the analysis of the long-run trends in structural transformation that we consider here.

expected length of their future career. As a result workers with a high idiosyncratic opportunity in a declining industry are more inclined to switch industries as they grow older, contrary to the data. In order to prevent these switches, retraining costs need to be increasing in age to offset the decline in the option value of being in a growing industry.

The third insight is that most of the impact of retraining costs is the on the trends in relative wages across sectors and not on the shares of workers employed. This is due to the aggregate technology being parameterized as near-Leontief to be consistent with historical trends in value added shares and relative prices in the U.S. (Ngai and Pissarides, 2008a).

The final insight is that because more workers will line up in the service sector when structural transformation accelerates unexpectedly, such an acceleration reduces wages in the service sector in the decades directly after. This reduction disproportionately affects the career earnings outlook of older workers in the service sector when the shock hits. In the longer-run the shock reduces relative wages in agriculture and manufacturing.

2.2 Cohort Effects and Structural Transformation

In this section we document stylized facts of worker reallocation across industries that motivate our model in Section 2.3. We focus our attention on a classic three-industry view of the U.S. time series, with some additional results from a large international sample presented for comparison. Details of the data construction and results from alternative industry decompositions are reserved to the appendix.

Our baseline analysis uses the United States census microdata spanning 1870–2010, taken from IPUMS (Ruggles *et al.*, 2010). We study the structural transformation of employment, which is constructed using the reported industry of employed

workers with valid responses. IPUMS has devoted substantial effort to harmonizing responses to these and other key variables over time and across countries. We aggregate detailed industry classifications to the standard three broad industry groups: agriculture, manufacturing, and services.⁶ We impose no other sample restrictions, because we want the results derived from microdata to be consistent with aggregate trends.

Our main empirical finding is that much of structural transformation can be accounted for by new workers who enter growing sectors disproportionately. Figure 2.1 provides a visual representation of this finding for the United States. It plots sectoral employment shares against time, with each individual line representing a distinct decade-of-birth cohort followed over their working life. Ignoring for a moment the distinct lines, the general employment patterns are clear: the decline of agriculture; the inverse-U shape of manufacturing; and the rise of services. The individual lines show that within a particular year younger cohorts had different employment patterns than older ones. For example, in 1900 the younger cohorts had about 15 percent lower employment shares in agriculture and correspondingly higher employment shares in manufacturing and services. The between-cohort gaps within a year provide visual evidence of the importance of cohort effects for structural transformation.

To document this pattern more carefully we utilize a within-between accounting decomposition. Denote by $e_t(i)$ sector i 's employment share at time t and by $\Delta e_t(i)$ the change in sector i 's employment share between $t - 1$ and t . We decompose this total change into two pieces: the portion accounted for by changes in the employment share of the cohort who is age h at time t , $n_t(h)$; and the portion that is accounted for by the changes in the employment patterns of each cohort, $e_t(i; h)$. The usual

⁶Agriculture includes all of farming, forestry, and fishing. Manufacturing includes also mining and construction. Services includes utilities.

decomposition holds:

$$\underbrace{\Delta e_t(i)}_{\text{total}} = \underbrace{\sum_h \overline{n_t(h)} \Delta e_t(i; h)}_{\text{within-cohort}} + \underbrace{\sum_h \Delta n_t(h) \overline{e_t(i; h)}}_{\text{between-cohort}}, \quad (2.1)$$

where Δ denotes differences between $t - 1$ and t and bar denotes averages between $t - 1$ and t .

We use variance-covariance accounting to perform the decomposition. The within and between shares are simply the covariance of the within and between terms with the total, relative to the variance of the total. This accounting procedure is identical to classical ANOVA. It can also be implemented in a straightforward way by taking the estimated coefficients from regressing the within and between components on the total component without a constant.

Table 2.1: Between Cohort Share of Structural Transformation

Sample	Total	By Industry		
		Agriculture	Manufacturing	Services
United States	53%	71%	36%	41%
International	53%	66%	23%	48%

Table 2.1 shows the result of this decomposition. In the first row we show the results for the United States, where we study reallocation between each consecutive pair of censuses, usually taken a decade apart. In the first column we show the results from pooling all three industries. In this case, the between cohort share of structural transformation is just over half, meaning that a little more than half of structural transformation is accounted for by the propensity of new cohorts of workers to work in growing sectors. The remaining columns show the results separately for agriculture, manufacturing, and services. The between share is highest for agriculture and

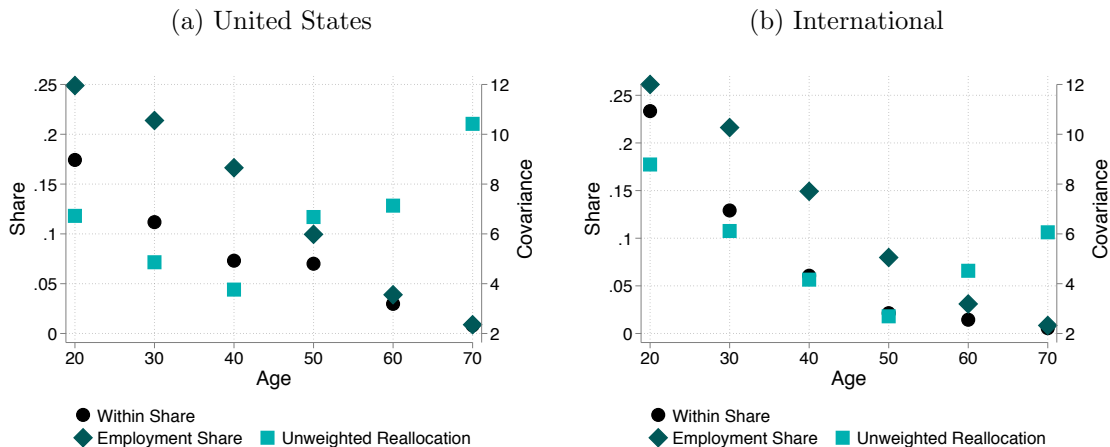
somewhat lower for manufacturing and services. These findings are consistent with the work of Kim and Topel (1995), who showed that between-cohort reallocation was central to the decline of agriculture in South Korea, and support the focus of Porzio and Santangelo (2017) on the role of cohorts for agriculture versus non-agriculture.

Although we focus on the United States, the underlying patterns are quite similar for the international sample, which includes 201 nationally representative surveys from 59 other countries, allowing us to decompose structural transformation across 142 consecutive survey pairs around the world. The results are shown in the second row of Table 2.1. The overall share of 53 percent is almost identical to the share for the United States. The shares by industry are also quite similar, with again a much larger role of the between share in agriculture.

Just over half of structural transformation is driven by the replacement of old cohorts by new ones. Further, much of the within-cohort reallocation happens early in the life cycle. To document this point, we exploit the fact that our accounting equation is additive in age. We then decompose the within share into the portion that happens within a cohort for those aged 20–29 at time $t - 1$; those aged 30–39 at time $t - 1$; and so on. The results are shown as solid circles for the United States and the international sample in Figure 2.2. A further 15–25 percent of all structural transformation happens from the 20s, with the pace of reallocation slowing with age, somewhat more rapidly for the international sample.

Examination of the within component in equation (2.1) shows that it can decline for two reasons: because the employment share of a cohort falls with age (falling $\overline{n_t(h)}$); or because cohorts are less likely to switch sectors as they age (falling $\Delta e_t(i; h)$). It is useful for our purposes to distinguish between the effects of the employment share versus the unweighted reallocation. The diamonds in Figure 2.2 plot the average employment share by age. The squares, plotted against the right axis,

Figure 2.2: Within cohort effects by age



show the unweighted reallocation, which is the result of doing the same variance-covariance decomposition using only $\Delta e_t(i; h)$ as our measure of within. Although this no longer decomposes total reallocation, it does isolate the pure behavioral response. Indeed, we can see that both for the United States and the international sample the declining within share is driven primarily by a falling employment share by age. The unweighted reallocation effect is mixed: it falls in importance until age 40 or 50 before rebounding and becoming more important at older ages.

To summarize, the main contribution of our empirical work is to show that half of structural transformation seems to be accounted for by the fact that new cohorts disproportionately enter growing sectors. Much of the rest happens early in the life cycle, although this is driven more by demographics than by the behavioral responses of workers. These facts motivate us to write down a model of structural change where demographics and the employment choices of new workers play the central role.

One possible concern with our approach is that our between-cohort effects may proxy for other slow-moving trends that are the fundamental driving forces of struc-

tural transformation. For example, recent work has stressed the role of education and female labor force participation for structural transformation (Caselli and Coleman, 2001b; Rendall, 2017; Buera *et al.*, 2017; Ngai and Petrongolo, 2017). We test the importance of these factors by examining the share of structural transformation that happens within and between *gender* \times *marital status* \times *education* groups. We use binary gender and marital status categories and four education bins, which produce a total of 16 possible cells. Table 2.2 shows the corresponding accounting results. 16 percent of structural transformation happens between demographic cells, which is a much lower share than our cohort results above.⁷ This finding suggests to us that cohort is not simply a proxy for trends in education, female labor force participation, and so on. This suggests that new cohorts inherently account for much of structural transformation. We turn now to a model that captures this idea.

Table 2.2: Between Demographic Group Share of Structural Transformation

	Total	By Industry		
		Agriculture	Manufacturing	Services
United States	16%	10%	19%	18%

2.3 Structural Transformation with Career Choices

Our empirical findings suggest that new cohorts play a central role in the process of structural transformation. We now formulate a heterogeneous agent overlapping generations model of life-cycle career choice under switching costs and integrate it into a canonical model of structural transformation. In the next section we use the model to infer the nature of the adjustment costs and to perform several counterfactual

⁷Hendricks (2010) documents a similar facts for more detailed educational categories.

exercises that highlight how structural transformation is a race between demographics and technology.

2.3.1 Households

Demographics and cohorts

Because this paper is about the interaction between structural transformation and demographics, we start by defining the demographic structure of our model economy. The economy consists of a unit measure of households, that are made up of members indexed by age $h = 0, \dots, H$. Each year, a new cohort of size $N_t(0)$ is born into each household. The growth rate of these new cohorts is $n > 0$; cohorts aged H die with certainty, and younger ones die with probability $0 \leq \delta < 1$. The resulting law of motion for cohort size by age is given by:

$$N_t(h) = \begin{cases} (1+n)N_{t-1}(0) & h = 0 \\ (1-\delta)N_{t-1}(h-1) & h = 1, \dots, H \end{cases}. \quad (2.2)$$

The total size of the household (equivalently, total size of the population) is given by:

$$N_t = \sum_{h=0}^H N_t(h). \quad (2.3)$$

It also grows at rate n .

Preferences, consumption, and labor supply

The members of the household pool their income risk and maximize the present discounted value of the household's log consumption flow. The factor at which future consumption is discounted is β and this present discounted value equals

$$\sum_{t=0}^{\infty} \beta^t \ln C_t. \quad (2.4)$$

Here, following Ngai and Pissarides (2007a), the aggregate consumption level C_t is a CES aggregate of consumption $C_{a,t}$, $C_{m,t}$ and $C_{s,t}$ from the agriculture, manufacturing, and services industries, with C_t given by:

$$C_t = \left(\sum_{i \in \{a, m, s\}} \lambda_i C_{i,t}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}, \text{ where } \varepsilon < 1. \quad (2.5)$$

Here ε is the elasticity of substitution between the goods and services produced by the three main sectors. It determines how quickly households change their consumption patterns in response to trends in relative prices between sectors due to structural transformation; Ngai and Pissarides (2007a) show that $\varepsilon < 1$ generates trends in expenditure shares consistent with the data. The preference weights satisfy $\lambda_a + \lambda_m + \lambda_s = 1$.

We use log preferences here such that the real interest rate implied by the household's intertemporal choice, $r_t = \frac{1}{\beta} \frac{C_{t+1}}{C_t} - 1$, does not depend on population growth. Therefore, household's intertemporal choices are not affected by demographic factors. This allows us to isolate the effect of demographics on the transitional dynamics related to the (re-)allocation and training of workers that is the result of structural transformation.

Let $p_{i,t}$ for $i \in \{a, m, s\}$ be the price of goods and services of sector i , expressed in terms of units of the consumption aggregate C_t , which we use as our numeraire good throughout. The demand for each type of good $i \in \{a, m, s\}$ implied by the CES preferences is

$$C_{i,t} = \lambda_i^\varepsilon \left(\frac{1}{p_{i,t}} \right)^\varepsilon C_t. \quad (2.6)$$

The associated expenditure shares can be written as

$$s_{i,t} = \frac{p_{i,t} C_{i,t}}{C_t} = \lambda_i^\varepsilon p_{i,t}^{1-\varepsilon}. \quad (2.7)$$

The assumption that $\varepsilon < 1$ is sufficient to ensure that $s_{i,t}$ is increasing in $p_{i,t}$, which is consistent with cross-country evidence (Ngai and Pissarides, 2007a).

Households do not incur any disutility from working and so supply their labor inelastically (as in Ngai and Pissarides, 2007a; Herrendorf *et al.*, 2014b). We deviate from this existing literature in allowing workers to make career choices subject to training costs and retraining costs to switching between sectors. Since the introduction of these training frictions in the labor market is the main contribution of this paper, we present them in a separate subsection below. Before that, however, it is useful to first consider the firms' decisions that determine the supply side and labor demand schedules of our economy.

2.3.2 Firms

On the supply side of this economy, firms use labor as the only production factor. The production technologies in each of the sectors $i \in \{a, m, s\}$ are linear. We denote output of each respective sector by $Y_{i,t}$ and the sectoral Total Factor Productivity (TFP) by $A_{i,t}$, such that

$$Y_{i,t} = A_{i,t}L_{i,t}. \quad (2.8)$$

where $L_{i,t}$ is the amount of labor used in production in sector i . Note that, because workers differ in their productivity levels in this economy, $L_{i,t}$ is measured in terms of efficiency units of labor. What makes this a model of structural transformation is that we assume that the three sectors in the economy are subject to three different rates of TFP growth, g_i , such that

$$A_{i,t} = (1 + g_i) A_{i,t-1}. \quad (2.9)$$

Each of the three sectors is perfectly competitive in that firms are price and wage takers, and that there is free entry. Free entry of firms occurs until price equals the

average (and marginal) cost of production:

$$p_{i,t} = \frac{w_{i,t}}{A_{i,t}}. \quad (2.10)$$

Here $w_{i,t}$ is the real wage paid per efficiency unit of labor in sector $i \in \{a, m, s\}$ in period t . An important difference between our model and other studies of structural transformation is that we consider sector-specific labor markets. This is the reason that $w_{i,t}$ is not equated across sectors and, hence, it is denoted by a subscript i .

2.3.3 Career Decisions and the Labor Supply

The reason that wages differ between sector-specific labor markets is that individual workers' labor supply is not perfectly elastic across sectors. That is, individual workers do not simply choose to work for the sector that pays the highest wage. Instead, their sectoral choice is affected by three particular factors.

First, each worker receives an idiosyncratic sector-specific productivity shock $z_{i,t}$ in each period. These shocks are drawn from an exponential distribution with mean 1. For notational purposes, we combine these three shocks in the vector, $\mathbf{z}_t = [z_{a,t}, z_{m,t}, z_{s,t}]'$.

Second, it is costly for individual workers to get *trained* to acquire the skills necessary to work in a particular sector at the beginning of their career at age $h = 0$. Finally, it is also costly for them to get *retrained* in case they decide to work in a different sector mid-career at age $h > 0$. These latter two factors, i.e. the *training* and *retraining* costs, are the labor market frictions that are the main focus and contribution of this paper.

Both the *training* and *retraining* costs reflect that it takes workers time to initially get trained to start their career in a particular sector and then to get retrained in case they switch sectors. We capture these costs in terms of two parameters. The

training-cost parameter $\phi \in [0, 1]$ is the fraction of the period when the worker is of age $h = 0$ that the worker spends on getting trained to work in a particular sector. The *retraining*-cost parameter $\gamma_h \in [0, 1]$ is age-specific and reflects the fraction of a period that a worker spends on being retrained when he or she decides to switch sectors of employment after the initial training at age $h = 0$.

Household members share their income and are fully insured against these costs. Because of this, each individual worker chooses his or her career path to maximize the expected present discounted value of lifetime earnings. At time t , this choice depends on the worker's age, h , industry of employment, i , and productivity shocks \mathbf{z}_t . The expected present discounted value of net future lifetime earnings by the individual equals $V_t(i, h; \mathbf{z}_t)$.

Given that the workers make optimal career decisions to maximize their $V_t(i, h; \mathbf{z}_t)$, we can write the expected present discounted value of lifetime earnings as the following Bellman equations. At age $h = 0$ the worker is not employed in a particular sector yet. She chooses an initial sector to start her career, taking into account the productivity shocks \mathbf{z}_t and the fact that in order to get trained she will only work a fraction $(1 - \phi)$ of the first period of her career. The Bellman equation associated with this choice reads

$$V_t(0; \mathbf{z}_t) = \max_{i \in \{a, m, s\}} \left\{ (1 - \phi) z_{i,t} w_{i,t} + \frac{1 - \delta}{1 + r_t} \mathbb{E}_t V_{t+1}(i, 1; \mathbf{z}_{t+1}) \right\}. \quad (2.11)$$

Here r_t is the real interest rate in period t and \mathbb{E}_t is the expectation conditional on information available at time t . This expectation is over all possible realizations of the worker-sector-specific productivity shocks, i.e. \mathbf{z}_{t+1} . The value function on the left-hand side does not have a sector index here because workers are not yet employed in a specific sector at the beginning of their career.

At age $h > 0$ the worker has started a career and the Bellman equation that

determines the value for a worker of age h employed in industry i in period t and faced with productivity shocks \mathbf{z}_t is given by

$$V_t(i, h; \mathbf{z}_t) = \begin{cases} \max_{j \in \{a, m, s\}} \{(1 - \mathbb{I}(j \neq i) \gamma_h) z_{j,t} w_{j,t} + \frac{1-\delta}{1+r_t} \mathbb{E}_t V_{t+1}(j, h+1; \mathbf{z}_{t+1})\} & \text{if } h = 1, \dots, H-1 \\ \max_{j \in \{a, m, s\}} \{(1 - \mathbb{I}(j \neq i) \gamma_h) z_{j,t} w_{j,t}\} & \text{if } h = H \end{cases} \quad (2.12)$$

Here, the indicator function $\mathbb{I}(j \neq i)$ reflects that a worker spends a fraction γ_h of her time on being retrained in the period when she decides to switch sectors and that she does not have to spend any time on retraining if she remains in the same sector. The latter case reflects that workers of age $h = H$ die with certainty and, thus, do not have a continuation value to their careers.

The result is that workers' career decisions involve a dynamic discrete choice problem. This discrete choice problem can be summarized in four variables that are important for the equilibrium dynamics of the labor supply and, thus, the economy.

The first two of these variables have to do with workers' *training* decisions at the beginning of their career, at age $h = 0$. First is the probability that a worker of age $h = 0$ is trained to work in sector i at time t , $\Phi_t(i)$. Second is the average number of efficiency units of labor that these workers supply to sector i at time t , $\tilde{z}_t(i; 0)$. The zero here denotes the age of the workers making the training decision.

The latter two variables are determined by the *retraining* decisions in period t , which depend on the worker's age h , the industry of employment i , and the productivity shocks, \mathbf{z}_t . The first is the probability that a worker of age $h > 0$ who works in sector i at time t decides to get retrained and starts working in sector j , $\Gamma_t(i, j; h)$. The other variable is the average productivity level for working in sector j of workers of age h who switch from sector i to j in period t , $\tilde{z}_t(i, j; h)$.

The assumption that the idiosyncratic sector-worker-time specific productivity shocks have exponential distributions allows us to solve $\Phi_t(i)$, $\tilde{z}_t(i; 0)$, $\Gamma_t(i, j; h)$, and $\tilde{z}_t(i, j; h)$ in closed form as a function of the wages in each of the sectors and the continuation values of working in each of the sectors. These closed-form solutions are algebraically intense and we therefore leave them for Section B.4 in the Appendix. What is important in the rest of our analysis is that these four variables are sufficient to describe the dynamic evolution of the labor supply in our model.

2.3.4 Equilibrium

Product markets

Because output is only used for consumption, product market equilibrium requires

$$Y_{i,t} = C_{i,t}, \text{ for } i \in \{a, m, s\}. \quad (2.13)$$

Through the inverse demand function implied by (2.6), this determines the relative prices as a function of relative demands as

$$p_{i,t} = \lambda_i \left(\frac{C_t}{C_{i,t}} \right)^{\frac{1}{\varepsilon}} = \lambda_i \left(\frac{Y_t}{Y_{i,t}} \right)^{\frac{1}{\varepsilon}}. \quad (2.14)$$

where we define aggregate output as $Y_t = C_t$.

Labor markets

Free entry of producers drives down the price to equal the average cost of production, which, using (2.10) and (2.14), determines real wages as a function of relative output levels, $Y_{i,t}$, and TFP levels, $A_{i,t}$:

$$w_{i,t} = A_{i,t} p_{i,t} = A_{i,t} \lambda_i \left(\frac{Y_t}{Y_{i,t}} \right)^{\frac{1}{\varepsilon}}. \quad (2.15)$$

Using the production functions, (2.8), we can write this expression for the real wages in terms of labor inputs and relative productivity levels

$$w_{i,t} = \lambda_i A_{i,t}^{\frac{\varepsilon-1}{\varepsilon}} \frac{\left(\sum_j \lambda_j A_{j,t}^{\frac{\varepsilon-1}{\varepsilon}} L_{j,t}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{1}{\varepsilon-1}}}{L_{i,t}^{\frac{1}{\varepsilon}}}. \quad (2.16)$$

These are the inverse labor demand functions that determine how many efficiency units of labor firms will hire at given real wages and productivity levels.

If labor markets were frictionless and labor was homogenous then the labor supply for each sector would be perfectly elastic. As a result $w_{i,t} = w_t$, as in canonical models of structural transformation. Here the dynamics of the labor market are more complicated because the labor supply depends not only on current wages, but also on past and current career decisions of workers, which in turn depend on past and future wages.

In fact, because it is costly for workers to get trained to work in different sectors, the age-industry structure of the labor supply in this economy is a state variable whose law of motion is determined by demographics and the career decisions of workers. We derive the law of motion of the labor supply in three steps. In the first, we follow how many workers of age h work in sector i at time t . We denote this number by $E_t(i; h)$. In the second step, we consider how many efficiency units of labor these workers supply, based on their endogenous, career choices. In the final step, we aggregate these efficiency units over workers of all ages to get the aggregate labor supply for each sector i at time t .

The number of workers of age $h = 0$ who work in sector i at time t is equal to the number of persons of age $h = 0$ at time t , i.e. $N_t(0)$, times the fraction of them who decide to get trained to work in sector i , i.e. $\Phi_t(i)$. Thus,

$$E_t(i; 0) = \Phi_t(i) N_t(0). \quad (2.17)$$

The number of workers of age $h > 0$ who work in sector i at time t , $E_t(i; h)$ is equal to the number of workers of age $h - 1$ who worked in the sector a period ago and didn't die, $(1 - \delta)E_{t-1}(i; h - 1)$, times the share of them who do not switch sectors, $\Gamma_t(i, i; h)$, plus the sum of the workers of age $h - 1$ who worked in other sectors in period $t - 1$ that survived and decided to switch to sector i in period t . Mathematically, the boils down to

$$E_t(i; h) = \sum_{j \in \{a, m, s\}} (1 - \delta) \Gamma_t(j, i; h) E_{t-1}(j; h - 1), \text{ for } h = 1, \dots, H. \quad (2.18)$$

The $3 \times (H + 1)$ -dimensional tuple, $\{E_t(i; h)\}_{i, h}$ is the state variable in this economy that determines the labor supply.

This state variable is measured in terms of numbers of workers. The labor inputs for each sector, $L_{i, t}$, are measured in terms of efficiency units of labor instead. The labor supply can be transformed from number of workers into efficiency units of labor by multiplying the number of workers by their average productivity level that depends on their career choice and by the net (of training and retraining time) number of hours that these workers supply. To do this, we denote the number of efficiency units of labor supplied to sector i by workers of age h in period t by $L_t^s(i; h)$. This allows us to write

$$L_t^s(i; 0) = (1 - \phi) \tilde{z}_t(i; 0) \Phi_t(i) N_t(0), \quad (2.19)$$

and

$$L_t^s(i; h) = \sum_{j \in \{a, m, s\}} (1 - \mathbb{I}(j \neq i) \gamma) \tilde{z}_t(1 - \delta)(j, i; h) \Gamma_t(j, i; h) E_{t-1}(j; h - 1), \text{ for } h = 1, \dots, H. \quad (2.20)$$

These equations define the industry-age-specific labor supply curves, in terms of efficiency units of labor.

Equilibrium in the labor market is when the sector-specific real wages, $w_{i,t}$, adjust such that the total number of efficiency units of labor demanded in a sector, i.e. $L_{i,t}$, equals the aggregate supply of efficiency units of labor to this sector. That is,

$$L_{i,t} = \sum_{h=0,\dots,H} L_t^s(i; h). \quad (2.21)$$

Here, the left-hand side variable depends on the real wages through the inverse labor demand function, (2.16), while the right-hand side variables depend on the real wages through the workers' career choices.

2.4 The Impact of (Re-)Training Costs

In this section we consider the impact of (re-)training costs on the equilibrium dynamics of our model. We do so by comparing the dynamic equilibrium path of our economy with such costs with a baseline case in which such costs are not present. We call this baseline case the *Flexible Benchmark*. We illustrate the impact of (re-)training costs in four steps.

First, we describe the main properties of the equilibrium of the Flexible Benchmark case. This case is very similar to the model introduced in Barany and Siegel (2018) and we discuss the similarities as well as emphasize the properties that are important to understand when we add (re-)training costs for workers. The most important property of the flexible benchmark is relative wages in the service sector are increasing compared to those in manufacturing and agriculture.

Next, we show that retraining costs accelerate the process of structural transformation in the economy rather than slow it down. We do so with an example in which retraining costs are the same for workers of all ages (*flat retraining costs*). The counterintuitive result that retraining costs accelerate structural transformation is because there is an option value to working in the service sector in anticipation of fu-

ture wage gains. This option value is higher for young than for old workers. Because of this, under flat retraining costs the model has the counterfactual implication that older workers disproportionately switch back from the service sector to agriculture and manufacturing.

In the third step of this section we show that the absence of such career switchbacks in the data implies that, in the context of our model, retraining costs need to be increasing in age.

Finally, we explain why retraining costs in this model mainly affect the trends in relative wages across sectors rather than the trends in employment shares. This is a consequence of the near-Leontief preferences in the parameterization of our model.

The fact that sectoral employment shares are not affected much by retraining costs does not mean that these costs have not effect on output. We show that retraining costs reduce output in for two reasons. The first is that they siphon off labor from production into training. The second is that they distort the workers labor supply decisions reducing the efficiency units of labor employed in each sector.

2.4.1 Flexible Benchmark and Solution Method

Flexible Benchmark

Throughout the rest of this section we use the case in which there are no (re-)training costs, i.e. $\phi = \gamma_h = 0$ for $h = 1, \dots, H$, as our main baseline for comparison. This *flexible benchmark* is a useful baseline because it is similar to the transitional dynamics studied in other analyses of structural transformation. Most notably, our flexible benchmark is very similar to the equilibrium dynamics in Bárány and Siegel (2018).

When workers do not face any training or retraining costs, their period-by-period

labor supply decision in this model simply involves choosing to work in the sector i that pays the maximum compensation given their idiosyncratic productivity draws, $z_{i,t}w_{i,t}$. Thus, in this case, workers' career choices neither depend on their future career opportunities, nor on their initial industry of employment, nor on their age. Moreover, because we have abstracted from capital as a factor of production, the level of output per worker is also not affected by the population growth rate and workers' life expectancy. This means that we use the this baseline to consider how workers' career decisions change relative to it when workers face (re-)training costs and how the aggregate level of output per worker is affected by these changes in career decisions.

Because we rely on numerical methods for our analysis, we have to choose a set of baseline parameters to evaluate the dynamics of the flexible benchmark. Following Ngai and Pissarides (2008a), we choose $\varepsilon = 0.1$, which is in the range of estimates implied by postwar U.S. national income data. We discipline our choice of the other parameters by having the flexible benchmark match the historical U.S. employment shares in agriculture, manufacturing, and services at the beginning and end of our sample, i.e. 1870 and 2010. We also match the average annualized historical growth rate of real GDP per capita over the sample. This calibration is described in more detail in Appendix B.5. We transform the annualized parameters in our model to reflect the length of a period which we set to 10 years.

The demographic parameters, i.e. the population growth rate, n , and the mortality rate, δ , do not matter for equilibrium in the flexible benchmark. Thus, we cannot use the flexible benchmark path to quantitatively discipline them. Instead, we choose n to match the average annual population growth rate in the U.S. between 1870 and 2010 and δ to match the average annual mortality rate for persons aged 10-70 born between 1904 and 1942 from Carter *et al.* (2006). In addition, we set the discount factor to $\beta = 0.95$ (annualized).

We assume that persons live for 6 periods in our model, i.e. $H = 5$. Given the period length of 10 years, one can interpret this as covering individuals from age 10 through 70 (similar to the data we analyzed).

The equilibrium path of our economy in the flexible benchmark closely resembles that of the model introduced in Bárány and Siegel (2018).⁸ The main difference is that the labor supply in our model is made up of cohorts of workers who make lifetime career decisions. These career decisions, and how they compare to the evidence we presented in Section 2.2 is what we focus on here.

In the flexible benchmark, the choice of the sector that pays the highest compensation does not depend on a worker's age. As a result, all cohorts make the same career decisions. Panels (i) - (a) through (c) from Figure 2.3 show this. They are the model-equivalent of Figure 2.1. Contrary to the data, in the flexible benchmark the fraction of workers that are employed in each sector is the same across cohorts. This is why the lines in the panels in row (i) of Figure 2.3 overlap.

This also shows that the *changes* in employment shares across sectors are the same across cohorts in the benchmark. Table 2.3 reports the between-cohort share from the ANOVA of the changes of aggregate employment shares for three model specifications. These are the model-equivalent of Table 2.1, and the first row shows it for the flexible benchmark. These between-cohort shares are much lower than in the data. Even though all cohorts make the same career decisions, the between share is not zero. This is because average employment patterns over their life cycle differ across cohorts because they are alive during different periods.

Panel (a) of row (i) of Figure 2.4 shows the trends in the relative (log) wages across sectors that drive workers' career decisions. In the flexible benchmark wages in agriculture initially exceed those in manufacturing and services. Most importantly,

⁸We illustrate this using a detailed set of results in Appendix B.6

Table 2.3: ANOVA in Different Model Specifications

	Total	By Industry		
		Agriculture	Manufacturing	Services
		<i>(i) Flexible benchmark</i>		
		$\phi = 0, \gamma_h = \{0, 0, 0, 0, 0\}$		
Share Between Cohorts	18%	18%	18%	18%
		<i>(ii) Flat retraining costs</i>		
		$\phi = 0.65, \gamma_h = \{.5, .5, .5, .5, .5\}$		
Share Between Cohorts	53%	24%	66%	-5%
		<i>(iii) Increasing retraining costs</i>		
		$\phi = 0.65, \gamma_h = \{.5, .5, .5, .9, .999\}$		
Share Between Cohorts	20%	45%	23%	52%

relative wages in the service sector increase over time. This is driven by the increase in the relative price of services along the transitional path of this economy. As Ngai and Pissarides (2007a) point out, the complementarity of the goods and services produced by the three sectors in this economy and the relatively low productivity growth results in an increase in the relative price of services (and an increase in the share of value added) along the transitional path.

Two other effects reduce the trend in relative wages in services. The first is the low productivity growth in services, which puts downward pressure on real wage growth in the service sector. The second is the selection of workers into the service sector. Consistent with the evidence provided in Young (2014), workers of increasingly low average productivity are drawn into services. That is, \tilde{z}_s, t declines over time. These

workers are drawn into the service sector by the increase in relative wages driven by the different productivity growth rates and resulting relative price trends across sectors. This downward trend in labor quality in the service sector also puts downward pressure on the growth rate of average labor productivity (per worker).

On net, however, the productivity and worker-selection effects are smaller than the relative price effect. As a result, in this economy real wages in the sector with the lowest productivity growth grow the fastest.

Panel (b) of row (i) of Figure 2.4, for $h = 1, \dots, H$, shows the covariance between the within-cohort and aggregate changes in employment shares, normalized by the variance of the aggregate changes. This ratio, from (2.1), is the regression coefficient of a regression of the age-specific changes in the employment shares on the aggregate changes in the employment shares. Because the career profiles of each cohort change in lockstep with the aggregate distribution of employment, these regression coefficients are 1 for all cohorts in the flexible benchmark. This stands in stark contrast to the variation in the data we documented in Figure 2.2

Thus, compared to the data, the flexible benchmark generates much less between-cohort variation in career profiles. Moreover, in the absence of (re-)training costs, changes in each cohort's career profile follows that of the overall economy, which is not the case in the data.

Implementation of Extended Path method

Because workers' career choices in the flexible benchmark do not depend on their previous decisions or age, the flexible benchmark case does not have a state variable and can be solved relatively straightforwardly on a period-by-period basis. In the presence of retraining costs, however, the equilibrium path of this economy depends on the initial age-industry distribution of the labor supply, $\{E_0(i; h)\}_{i,h}$.

This equilibrium path can be reduced to a path of (real) wages in each of the sectors, $\{w_{i,t}\}_{i,t}$ that, at each point in time, equates the demand and supply in each of the three labor markets. When the wages result in equilibrium in the labor market, Walras' Law implies that the product market will be in equilibrium as well. Because the equilibrium depends on the complicated evolution of the age-industry distribution of the labor supply, which in turn is determined by the workers' dynamic discrete career choices, it is not possible to find a closed-form solution for the equilibrium path. Instead, we have to resort to numerical methods.

The solution method that we use, described in Section B.7 of the Appendix, is an application of the "Extended-Path" method, which was first discussed in Fair and Taylor (1983) and applied in, for example, Greenwood and Yorukoglu (1997) and Hobijn *et al.* (2006). Normally, the extended path method solves the transitional dynamics of a model between one steady state in period $t = 0$ and another at $t = T$. Because our model does not have a steady state or balanced growth path, we use a slightly different approach.

We solve the transitional dynamics of our model economy over the period from $t = -\tilde{t}_l$ through $t = T + \tilde{t}_r$. We assume that the initial state of the economy, at $t = -\tilde{t}_l$, is the one from our flexible benchmark. Moreover, we assume that the economy is on a balanced growth path, in which all sectors grow at the same rate, with no (re-)training costs after $t = T + \tilde{t}_r$. The reason that we add the left- and right-padding, i.e. $\tilde{t}_l > 0$ and $\tilde{t}_r > 0$, to the solution path is to reduce the impact of the assumed initial and final conditions on the part of the solution path, namely $t = 0, \dots, T$, that we focus on.

2.4.2 Retraining Costs Accelerate Structural Transformation

The first thing we illustrate is that the addition of retraining costs to the model *accelerates* rather than slows down the process of structural transformation, as captured by the shift in employment from agriculture through manufacturing to services.

We illustrate this property of the model for a case with flat retraining costs. In particular, we look at the case where $\phi = 0.65$ and $\gamma_h = 0.5$ for all $h = 1, \dots, H$. Under the restriction of flat retraining costs, i.e. $\gamma_h = \gamma$ for $h = 1, \dots, H$, this combination of parameters gets the closest to the between share for the United States reported in Table 2.1 and the cohort-career regression coefficients shown in Table 2.2.

Figure 2.5 illustrates the difference between the path of the sectoral employment shares in the flexible benchmark (hashed bars) and in the case with flat retraining costs. It shows that the employment shares of manufacturing and services in the early stages of the structural transformation are higher under the retraining costs than in the flexible benchmark. At first glance, this might seem like a very counterintuitive result, because we tend to think of adjustment costs slowing down adjustments rather than accelerating them.

The reason for this acceleration is that, when workers face retraining costs, career choices do not only depend on current (real) wages, $w_{i,t}$, but also on the career continuation values, $\frac{1}{1+r_t} \mathbb{E}_t V_{t+1}(i, 1; \mathbf{z}_{t+1})$. These continuation values reflect the option of being employed in a particular sector. This option value is particularly high in our model for the service sector, because it largely captures the present discounted value of the future increases in relative wages in the service sector *over the rest of a worker's career*. As a result, workers facing retraining costs choose to be employed in the service sector more than those that do not face retraining costs. They do so in

anticipation of future relative wage increases in services.⁹

In equilibrium, this increase in the labor supply in services results in higher employment in the service sector. It also subdues the trend in relative wages in the service sector compared to the flexible benchmark. This can be seen by comparing Panels (a) of rows (i) and (ii) of Figure 2.4. What happens is that retraining costs reduce the gross flows of workers between sectors that are largely driven by their idiosyncratic productivity levels. They, however, result in larger *net* flows of workers across sectors that are coordinated by the common career continuation values that the workers face.

The younger the worker, the higher this option value, and the more the worker's decision is driven by it. The result is that older workers put more weight on current wages and productivity shocks when they make their labor supply decisions. For example, an old worker in services that draws a high productivity shock, i.e. gets a good opportunity in manufacturing or agriculture, will switch back to one of the shrinking sectors with declining relative wages in the economy. This can be seen from the three panels in the second row of Figure 2.3. The left and middle panels show the cohort-specific employment shares in agriculture and manufacturing. The career switchbacks of older workers are reflected by the increases in these shares in the last two periods of each cohort's career. These increases are offset by a decline in the share of workers in services. Higher wage gaps between manufacturing (as well as agriculture) and services imply higher career option values. Higher wage gaps also make switchbacks more common. This is why their size increases in Figure 2.3 over the transition path.¹⁰

⁹Note that this result does not depend on our assumption of per-period independent idiosyncratic shocks the career continuation values also matter in case of persistent shocks.

¹⁰The career switchbacks are the result of our assumption that workers are subject to sector-specific productivity shocks each period. In our model these shocks are independent over time.

2.4.3 Retraining Costs Increase with Age

Comparing the three panels in the second row of Figure 2.3 with those in Figure 2.1, it is clear that the career switchbacks that the model generates are counterfactual. In order for them not to occur in our model, we need to assume that retraining costs are higher for older workers. These costs need to increase so that they offset the decline in the career option value of workers in shrinking sectors.

To illustrate the effect of increasing retraining costs over the life cycle, we increase the retraining costs for workers of age $h = 4$ and $h = 5$ to $\gamma_4 = 0.9$ and $\gamma_5 = 0.999$. This makes it very costly for older workers to change their sectors of employment, and greatly reduces career switchbacks. We call this the *increasing retraining costs* case and its equilibrium path is plotted in Row (iii) of Figures 2.3 and 2.4. This is an extreme example. However, it shows that with late-career high retraining costs, career option values play a larger role. Workers move more rapidly from manufacturing to services in anticipation of future relative wage increases. This is, at least qualitatively, more consistent with the cohort career employment patterns depicted in Figure 2.1. Our data of consecutive cross-sections does not allow us to consider individual employment transitions by workers, but evidence in Menzio *et al.* (2016) suggests that these transitions are, indeed, declining over the life cycle.

Note, however, that with more workers lining up in services the equilibrium trend of relative wages is lower, compared to the flexible benchmark and flat retraining costs cases. The left panel of Row (i) of Figure 2.4 shows this. It is most evident from the real wages in agriculture: in the flexible benchmark these were declining after 40 years, and with increasing retraining costs, these continue to increase. In

If they weren't, highly correlated shocks over time or across sectors would imply smaller career switchbacks.

fact, the difference in real wages in agriculture at the end of the quarter millennium that we consider is 200 log points. That is about a 700 percent difference. In fact, most of the impact of retraining costs is not on the employment shares but on the relative wages across sectors.

2.4.4 Bulk of Impact of Retraining Costs Through Relative Wages

The reason that most of the adjustment to retraining costs in this economy goes through relative wages is that preferences are almost Leontief. As is commonly done in studies of structural transformation (e.g. Ngai and Pissarides, 2008a), we have chosen a very low elasticity of substitution between the goods and services produced by the sectors in the economy, i.e. $\varepsilon = 0.1$. This elasticity is consistent with the trends in relative prices and value added shares of the three sectors in U.S. time series.

It implies, however, that the marginal rates of substitution between $C_{i,t}$ for $i \in \{a, m, s\}$ vary a lot in response to small changes in the quantities. These marginal rates of substitution affect the relative wages in this economy because they determine the elasticities of the sector-specific labor demand curves. The smaller ε the less elastic the labor demand curves. The inelastic labor demand curves mean that changes in the labor supply mainly result in changes in relative wages rather than changes in sectoral employment shares.¹¹

Figure 2.6 illustrates this for period $t = 50$. It shows the sector-specific labor supply and labor demand curves under the *flexible benchmark* (dashed) and frictional *increasing retraining costs* (solid). The downward sloping labor demand curves are

¹¹In the limiting case of Leontief preferences in which $\varepsilon \downarrow 0$ labor demand (in efficiency units) is pinned down completely by technology parameters and all the effect of retraining costs on the labor supply is reflected in changes in relative wages rather than employment shares.

very inelastic (close to vertical) and do not move much due to the changes in the wages in the other sectors from the flexible to the frictional case. The result is that the equilibrium employment shares do not change a lot between the flexible and frictional case. The shifts in the labor supply curves induced by the retraining costs therefore translate mostly into changes in the equilibrium real wages.

2.4.5 *Output Losses Due to Retraining and Allocation of Labor*

The fact that retraining costs do not affect sectoral employment shares much does not mean, however, that these costs do not affect the level of output in the economy. Output in the case of increasing retraining costs is much lower under increasing retraining costs than in the flexible benchmark for two reasons. The first is the loss in effective labor input due to the (re-)training time of workers. The dashed purple line in Figure 2.7 shows the size of the gap in output between the increasing costs and flexible benchmark cases due to (re-)training in log points over the transition path.

At the beginning of the transition path, the (re-)training costs result in a loss of more than a third of output. In large part, this is due to the training costs, ϕ , that all workers incur when they start their career. In addition, this reflects the retraining time of workers that switch sectors.¹² Note that this loss declines as more workers are employed in services and fewer gross flows of workers occur between sectors. At the end of the transition path, training time absorbs about a quarter of the time available from all workers, which is mainly due to the training cost of workers in the

¹²The level of these costs in this model is relatively high because the variance of the Exponential distribution of workers' idiosyncratic productivity shocks is one. That is, the size of this loss is partly driven by our distributional assumption about workers' productivity shocks that keeps the discrete choice problem of workers tractable and solvable.

first period of their lives, i.e. because $\phi = 0.65$.

The second reason for the output loss has to do with the allocation of workers across sectors. (Re-)training costs can have a worker choosing to go into services when her current productivity (and wage) is higher in manufacturing or agriculture. This is due to the option value of going into services. The solid line, labeled “Total” in Figure 2.7 adds this effect on output on top of that for the training time to show the total output loss in the increasing retraining costs case compared to the flexible benchmark.

The part due to the allocation of labor, i.e. the difference between the two lines, peaks after about 30 years at around 15 percent and then declines as the employment distribution shifts towards services. Thus, because (re-)training costs affect the allocation of workers across sectors they reduce average labor productivity in this economy. This effect eventually goes to zero when all workers are employed in services.

In sum, in the long run retraining costs result in about a 20 percent reduction in output, solely due to the training of young workers to prepare them for careers in the service sector. However, during the process of structural transformation, the retraining time involved with the reallocation of workers across as well as workers’ forward-looking career decisions double this impact. That is, the total output loss peaks after 30 years at around 40 percent of the level of output in the flexible benchmark.

2.4.6 Unanticipated Acceleration of Structural Transformation

The insight that most of the impact of retraining costs in a canonical model of structural transformation is on wages instead of employment shares is important also for understanding the effect of an unanticipated acceleration in structural transforma-

tion. We illustrate this using a scenario in which the economy is on the same path as in the previous subsections for *flexible benchmark* and *the increasing retraining costs* cases. At time $t = 0$ structural transformation unexpectedly accelerates in that both g_a and g_m permanently increase by 50 percent compared to their calibrated values.¹³

Figure 2.8 plots the changes in the paths of the employment shares and the log of the average real wage paid per employee by sector. The two panels in Row (i) of the figure show the results for the *flexible benchmark* and the panels in Row (ii) for the *increasing retraining costs*.

The left panels in each of these rows show the change in the employment shares by sector over the post-shock equilibrium path of the economy. Both panels show how an increase in g_a and g_m results in a faster reallocation of labor from the agriculture and manufacturing sectors to services. For the reasons we already explained in subsection 2.4.4 above, the impact of the acceleration in technological change on the employment transition patterns is very similar for both cases.

The impact of the shock differs mostly in terms of the paths of the sectoral wages in the two cases. This can be seen by comparing the right-hand side panels in Rows (i) and (ii) of Figure 2.8. To understand the impact of the shock on wages under retraining costs it is important to distinguish between the initial impact and the subsequent dynamics.

In the absence of the retraining costs there is no initial impact of the shock on wages because it does not affect the sectoral productivity levels, $A_{i,0}$ for $i \in \{a, m, s\}$. The presence of retraining costs makes more younger workers line up in services in anticipation of future wage gains in that sector in response to the shock. This increases the labor supply in services at the time of the shock and lowers real wages

¹³That is, g_a increases from 0.045 (annualized) to 0.0675 and g_m from 0.020 to 0.030.

in the service sector.¹⁴ For young workers, this decline in the wage in the service sector at $t = 0$ is offset by subsequent increases in expected earnings later on in their careers. For older workers employed in the service sector at the time of the shock, this decline in the real wage is not offset by future wage increases and, in hindsight, some of these workers would have preferred to be employed in agriculture or manufacturing but are stuck in the service sector due to their high retraining costs when the shock occurs.

In the subsequent periods the shock results in a higher increase in the growth rate of real wages in the service sector than in the other two sectors. In this sense, it results in a stagnation of relative wage growth in agriculture and manufacturing. In the first 5 decades after the shock, this stagnation is more profound for the case of increasing retraining costs when the allocation of workers across sectors is still affected by the initial allocation at the time of the shock. In the longer-run retraining costs dampen the effect of the shock on relative real wage growth across sectors for the same reason that they dampen the trends in relative wages in the left-hand side panels in Figure 2.4. This can be seen from the very different scales on the vertical axes of the two right-hand side panels in Figure 2.8.

2.5 Conclusion

Using data on sectoral employment patterns by birth cohort from more than 50 countries around the world, we show that the bulk of the shift in the allocation of employment across major sectors of economic activity that is typical for structural transformation is because younger cohorts are disproportionately employed in growing sectors. We argue that the importance of between-cohort differences in sectoral

¹⁴Most of these workers turn out to be drawn out of the agriculture sector, which is where wages increase in response to the shock.

employment shares is indicative of workers incurring substantial retraining costs when they switch sectors.

To illustrate the aggregate implications of such retraining costs, we introduce a model of structural transformation with overlapping generations that face (re-)training costs when they make their labor supply decisions. On the aggregate level, we follow Ngai and Pissarides (2007a) and we model structural transformation as being driven by different levels of TFP growth across sectors. These sectors produce different types of output that are gross complements in aggregate CES preferences and drive demand patterns.

Our main theoretical contribution is modeling the career decisions of different cohorts as a discrete choice problem. In each period workers decide what sector they would like to work in based on their current productivity levels, the future wage prospects in each of the sectors, and the potential retraining costs they incur.

To match the between-cohort contribution to structural transformation in our model, we need to include substantial training costs for workers at the beginning of their careers and retraining costs when they switch sectors. We obtain four important insights from our model once we introduce (re-)training costs.

First of all, (re-)training costs *accelerate* the reallocation of labor to the growing service sector rather than slow it down. This is because workers decide to take jobs in the service sector in anticipation of future relative wage increases.

Secondly, because there is no option value to being in services for workers at the end of their careers, some of these move back to agriculture and manufacturing. These counterfactual career switchbacks in the model suggest that retraining costs are increasing in age and very high for old workers.

Thirdly, the strong gross complementarity of the goods and services produced by the three sectors in our economy results in the bulk of the impact of the retraining

costs being on trends in relative wages across sectors rather than trends in employment shares.

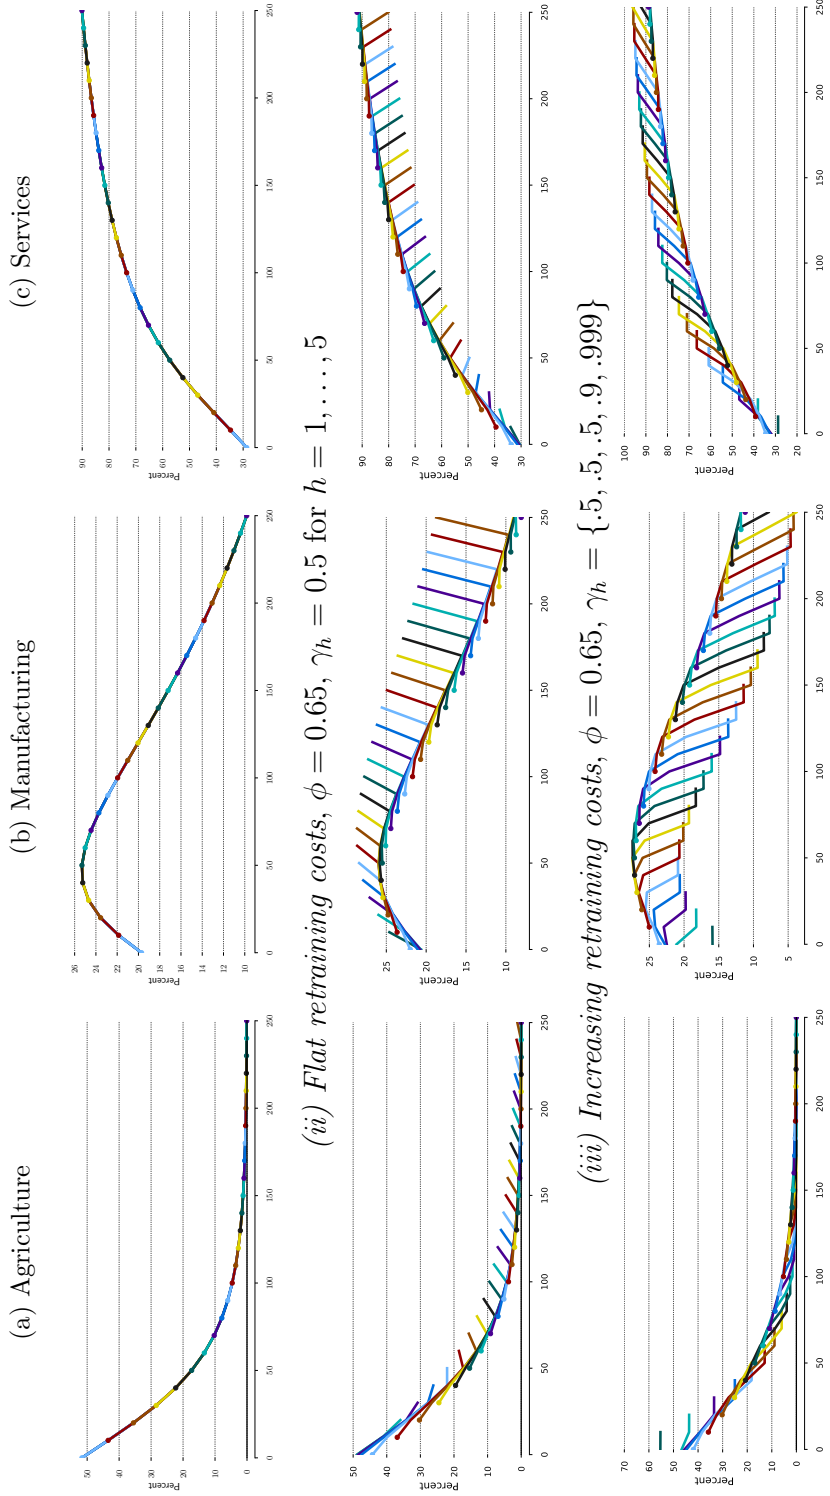
Finally, when structural transformation accelerates unexpectedly more young workers will choose to supply their labor in the service sector in anticipation of future relative wage growth in services. This reduces the service sector wage on impact of the shock which most negatively affects the career earnings outlook of older workers in the service sector. In the longer-run the shock reduces relative wages in agriculture and manufacturing.

Our main contribution here is the analysis of the evolution of the labor supply across cohorts in a general equilibrium framework. To isolate the effects of the assumptions we made about the labor supply, we deliberately embed our cohort-specific labor supply model in a simple model of structural transformation. This, deliberately stylized, framework reveals potentially large effects of labor market frictions, generally ignored in models of economic growth, on long-run economic outcomes.

The reallocation of labor due to structural transformation is only one of the potential applications of our theoretical contribution. For example, it can also be used to consider the cohort-specific effects of the labor market impacts of trade (Autor *et al.*, 2013, e.g.). As is common in models of economic growth, we have abstracted from workers' participation and retirement decisions. Adding those margins is a useful extension and the subject of future research.

Figure 2.3: Cohort-specific employment shares in three model specifications

(i) *Flexible benchmark*, $\phi = 0$, $\gamma_h = 0$ for $h = 1, \dots, 5$




Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010.


Figure 2.4: Wages and Cohort career profiles in three model specifications

(i) *Flexible benchmark*, $\phi = 0$, $\gamma_h = 0$ for $h = 1, \dots, 5$



(a) Real compensation per worker




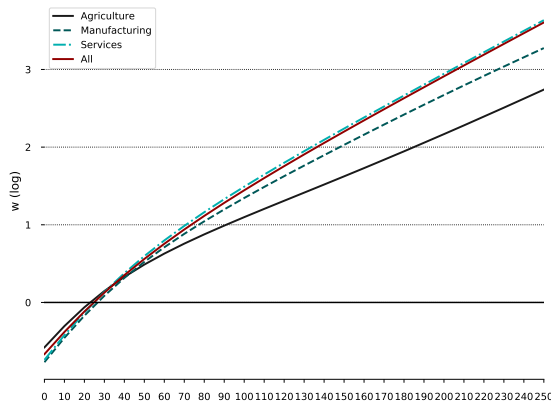
(b) Age-specific within regression coefficient



(ii) *Flat retraining costs*, $\phi = 0.65$, $\gamma_h = 0.5$ for $h = 1, \dots, 5$




(iii) *Increasing retraining costs*, $\phi = 0.65$, $\gamma_h = \{.5, .5, .5, .9, .999\}$



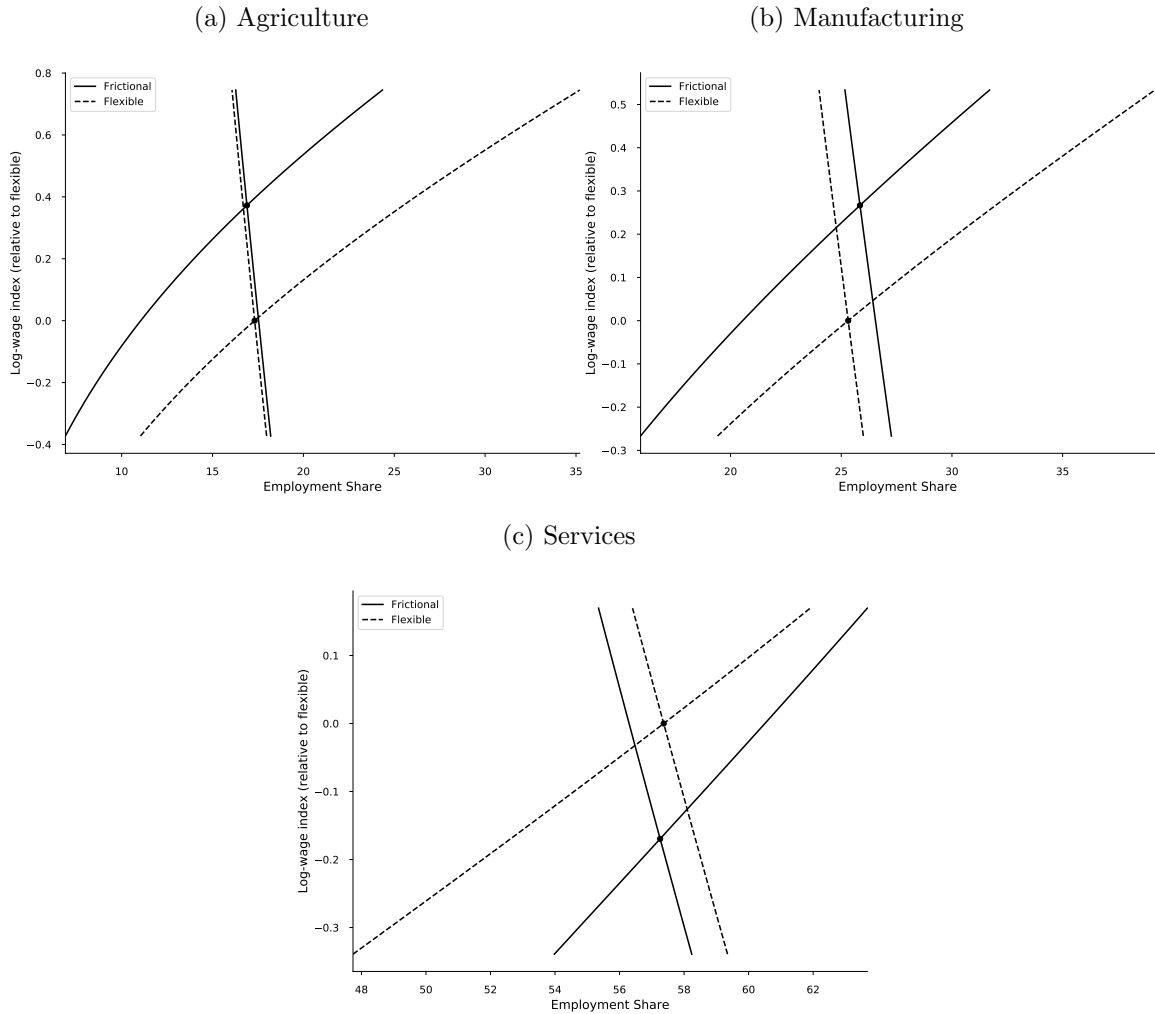
Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010. Age is measured in decades.

Figure 2.5: Sectoral employment shares in flexible benchmark and under flat retraining costs




Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010.

Figure 2.6: Sectoral labor supply and labor demand curves at $t = 50$



Note: Dashed lines are the labor demand (downward sloping) and labor supply (upward sloping) curves in *flexible benchmark*. Solid line are those in *increasing retraining costs* case. All curves are conditional on the equilibrium real wages in the other two sectors. The vertical axes show the log real wage in deviation from that in the flexible benchmark and the horizontal axes the employment shares $E_{i,t}/E_t$. The dots depict the equilibrium combinations of the log real wages and employment shares.

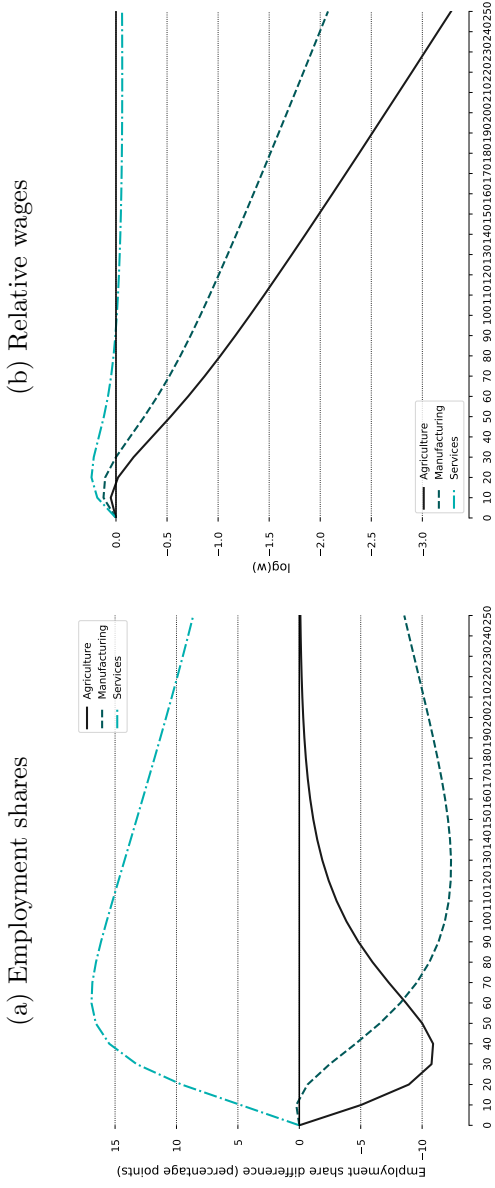
Figure 2.7: Difference in output levels between increasing retraining costs and flexible benchmark



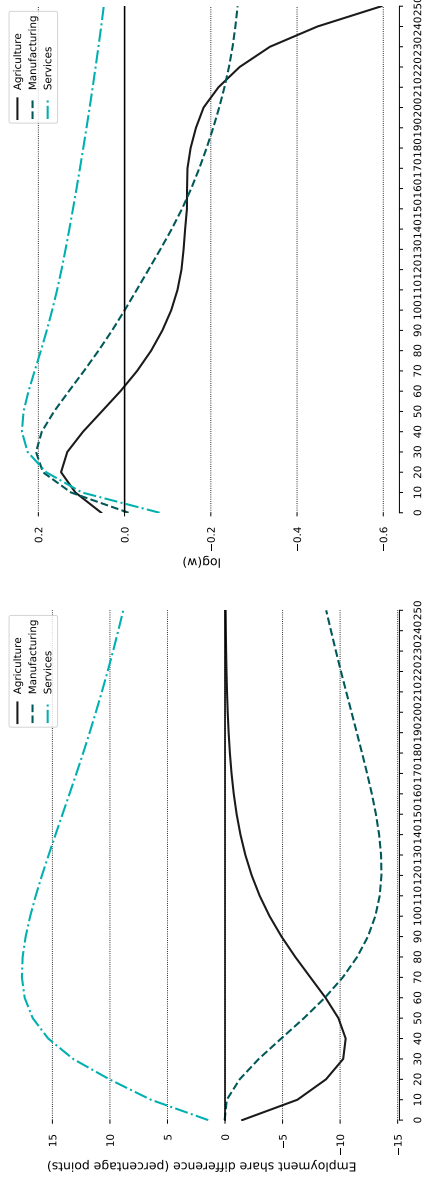
Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010. Gap is measured in terms of log points.

Figure 2.8: Impact of acceleration of structural transformation on employment and wages

(i) *Flexible benchmark*, $\phi = 0$, $\gamma_h = 0$ for $h = 1, \dots, 5$



(ii) *Increasing retraining costs*, $\phi = 0.65$, $\gamma_h = \{.5, .5, .5, .9, .999\}$



Note: The panels on the left show the percentage point difference between the sectoral employment shares with an acceleration in the rates of technological change in agriculture and manufacturing at time $t = 0$ and those where the rates of technological change are constant. The panels on the right show the differences in the log of the sectoral average wages paid per employee between the path with the shock and without it.

Chapter 3

WORLD POLARIZATION

3.1 Introduction

Job polarization is happening in the developed world. This is a fact that economists in labor and macroeconomics have documented.¹ In this paper, I present new stylized facts about the distribution of occupations in the world. With this, I expand current analyses of job polarization beyond a group of developed economies. I analyze a sample of 119 countries covering all levels of economic development to argue that job polarization is a global phenomenon.

Job polarization is concerned with the “hollowing out” of the employment distribution: occupations in the middle of the wage distribution decreasing their employment shares, while the shares for the occupations with higher and lower wages increase. These changes have been linked to the tasks that these occupations perform, resulting in three large groups: abstract, routine, and manual occupations. Abstract occupations mainly perform tasks requiring problem-solving, creativity, and persuasion typical of professional, managerial, and technical occupations. These occupations earn the highest wages, on average. Routine occupations perform tasks that follow well understood procedures, which makes them more susceptible to automation. Office clerks, and plant and machine operators are typical of this group of occupations, whose wages are in the middle of the wage distribution. The last group, manual occupations, perform tasks that require situational adaptability, visual recognition,

¹See, for instance, Autor and Dorn (2013); Beaudry *et al.* (2016); Cortés (2016); Goos *et al.* (2014).

and in-person interactions. Typical examples are occupations performing personal services, domestic and office cleaning, and construction and installation services.²

This polarization in labor markets has important implications for earnings, income distribution, and human capital. Job polarization is accompanied by wage polarization: wages in routine occupations have decreased compared to the other occupations (Firpo *et al.*, 2011). This polarizes the distribution of income, since routine occupations are in the middle of the wage distribution. Human capital has a strong occupation specific component (Kambourov and Manovskii, 2009). Therefore, changes in the occupational mix of the economy could lead to losses in accumulated human capital. These are all topics studied in the context of polarization. The scope of this paper, however, is on the employment distributions, and their evolution over time.

The two main stylized facts that I uncover are as follows. First, at any point in time, there is a strong link between a country's development level (measured by its real income per worker) and its occupational employment shares. This is what I call the *occupational development profile*. Second, and more importantly, this profile has shifted over time, resulting in world polarization: a lower routine development profile, coupled with higher manual and abstract development profiles. These facts are important because they portray how modern economic growth has affected the world distribution of occupations. This confirms that polarization is a global phenomenon.

Put differently, world polarization implies lower employment shares in routine occupations at *all* levels of development, increasing those in manual and abstract occupations. Take the cases of Spain and Peru: In 1985, Spain had a real income level comparable to Peru in 2014. Back then, Spain had a routine employment share of 38 percent, while Peru, 30 years later, had an employment share 11 percentage points

²For this paper, I borrow the classification from Cortés *et al.* (2014).

lower. The modern growth experience, then, is biased against routine occupations.

This finding contrasts to the results in the structural transformation literature. The link between income and employment shares is a stable one.³ Industries and occupations are closely related, so it would be tempting to argue that the link between development levels and occupational employment shares should also be stable. The data in this paper reveals otherwise.

Having a global perspective on polarization is useful in discerning among its causes. The polarization analyses so far have identified two main explanations: international trade and technical change (see, for example Goos and Manning (2007) & Autor and Dorn (2013)). International trade, and outsourcing routine employment from developed to developing countries would decrease the routine employment share in developed countries, while increasing it in developing countries. Polarization is happening at all levels of development, and also within industries. This means that an explanation based on technical change fits the patterns in the data better.

To analyze these development patterns, I follow the grouping principle for occupations and develop a polarization accounting framework at the *task* level. This serves two purposes. First, it allows to quantify the differences in task-specific productivities. This is important to determine the technological forces behind the occupational distributions. Second, it allows to study the process of technical change at a global level.

The main finding of this exercise is that technical change has been biased against routine occupations and biased in favor of manual occupations. Around the world, productivity growth has been highest in routine tasks, which makes workers to switch to different occupations. Productivity growth in abstract occupations has been the lowest, which goes in line with Baumol's cost disease (Baumol, 1967).

³See Herrendorf *et al.* (2014a) for a discussion.

These technological trends imply that as countries continue to develop, we can expect lower employment shares in routine occupations, and higher in manual occupations, worldwide. Through a vector autoregression (VAR) analysis, I forecast the path of productivity growth, and conclude that world polarization will continue. In the following years, then, the development profile for routine occupations will keep decreasing, and the profile for manual occupations will keep increasing.

This paper is organized as follows. The first part discusses how the world distribution of occupations changed between 1980 and 2014. The second one analyzes these changes through a polarization accounting framework, and forecasts its changes until 2050.

3.2 World Polarization Through the Lens of Occupational Development Profiles

To expand the sample of analyses of job polarization beyond the developed countries previously studied, I use two data sources with comparable occupational information: the International Labor Organization's ILOSTAT database, and census microdata from the IPUMS International Project. The result is an unbalanced panel of 119 countries, with an average period length of 19.4 years.⁴

The developed economies analyzed in Autor *et al.* (2003) and Goos *et al.* (2014) are relatively similar, which allowed for a simple comparison of employment shares without specifically taking into account their level of development. The countries I analyze are much more heterogeneous in their stages of development and occupational employment shares. In my broad sample it is thus important to study the link between the level of development and the occupational distribution of employment. To capture this link, I introduce the concept of the *occupational development profile* and use it

⁴The details of the data handling process are left to appendix C.1. All the results I report are qualitatively similar for unbalanced panel that I discuss here in the main text and a balanced panel.

to analyze polarization at the global level. The following subsection explains this concept in more detail.

Occupational Development Profiles

As countries develop, economic activity is reallocated across productive sectors (industries). This is one of the characteristics of modern economic growth, according to Kuznets (1973). More recent studies, like Duarte and Restuccia (2010) and Herrendorf *et al.* (2014a), confirm that this reallocation still holds for industries across various countries.

Kuznets' observations, however, also refer to changes in the occupational status of labor. This is the counterpart of this stylized fact for the distribution of employment across occupations. With this in mind, I refer to *occupational development profiles* as the link between development levels and occupational employment shares.

More precisely, I define $d_{j,t}$, the occupational development profile in occupation j and time t , as a function that relates a level of income y to its expected employment share in occupation j :

$$d_{j,t}(y) = \beta_{0,j,t} + \beta_{1,j,t}y + \beta_{2,j,t}y^2 \tag{3.1}$$

I estimate the parameters through ordinary least squares regressions, and table 3.1 presents the estimates for 1985 and 2014. Income levels correspond to the real, PPP adjusted GDP per worker relative to the U.S. level in 2000.⁵ This measure is intended to be comparable over countries and time. Since development profiles reflect cross-sectional data, these are the snapshots, at a point in time, linking the development process to its occupational employment shares. For the regression analyses, this is expressed in logarithmic terms, base 2.

⁵This reference point is the same throughout the analysis. This means that the U.S. series will only have a level of 1 in 2000.

Table 3.1: Occupational Development Profiles: 1985 & 2014

	1985			2014		
	Abstract	Routine	Manual	Abstract	Routine	Manual
Constant	33.984*** (3.103)	48.285*** (3.897)	17.731*** (6.089)	42.660*** (1.194)	29.434*** (0.831)	27.906*** (1.360)
Income	7.679*** (2.284)	4.369 (2.868)	-12.047** (4.481)	10.762*** (1.234)	-0.963 (0.859)	-9.799*** (1.406)
Income ²	0.417 (0.370)	-0.460 (0.465)	0.043 (0.727)	0.706*** (0.260)	-0.999*** (0.181)	0.294 (0.296)
N	29	29	29	119	119	119
R ²	0.773	0.802	0.831	0.716	0.565	0.800


OLS estimates of equation (3.1) for 1985 & 2014. The dependent variable is the occupational employment share, and income levels are included as the base 2 logarithm of PPP-adjusted per-worker GDP, relative to U.S. levels in 2000. Asterisks indicate statistical significance at 10 percent (*), 5 percent (**), and 1 percent (***).

Source: author's calculations using ILOSTAT, IPUMS & PWT.

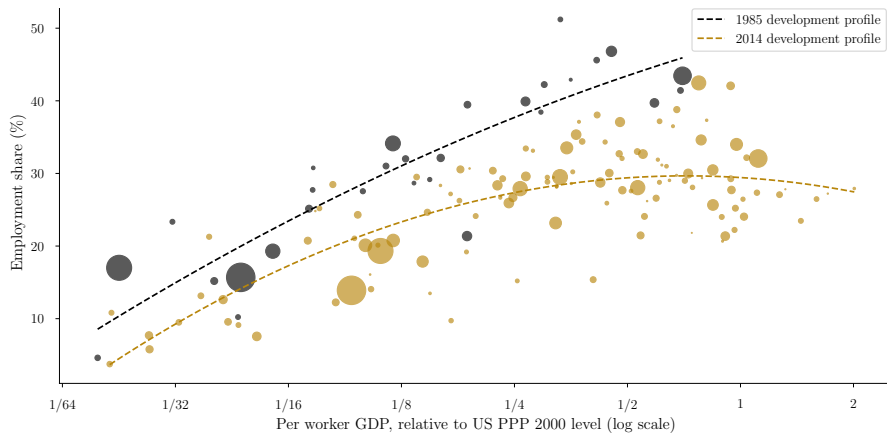
Figure 3.1 plots the estimated development profiles and employment shares for two years: 1985 and 2014. Each subpanel presents the information for an occupational group. The horizontal axes contain income levels, and the vertical axes occupational employment shares. Each circle represents a country, and its size is proportional to its level of employment. The information in black represents observations in 1985, and information in gold represents observations in 2014. It bears reminding that income is expressed relative to the U.S. level in 2000: during 1985, this income was 70 percent for the U.S., and 111 during 2014. This is to capture how occupational employment shares change during the different stages of development.

Figure 3.1: Occupational Development Profiles: 1985 & 2014

(a) Abstract occupations



(b) Routine Occupations



(c) Manual Occupations



Source: author's calculations using ILOSTAT, IPUMS & PWT.

This figure shows that there is a strong link between a country's income level, and its occupational employment shares. Typically, a country with a low income level also has a low share of its employment in abstract occupations. Countries with higher incomes, also have higher employment shares in abstract occupations. This is the case during 1985 and 2014. Routine occupations show the same pattern; countries with low levels of income have low employment shares, and tend to increase them as they develop. Finally, manual occupations show the opposite pattern. Low levels of income are associated with very high shares in manual occupations.

These development profiles capture the link between the distribution of employment across occupations and income levels at a point in time. For a given year, the movements *along* these curves tell us how the occupational employment shares are expected to change as countries develop. Take, for example, a country like Morocco. In 1985 its income level is 1/16 of the U.S. level in 2000, and its employment share in manual occupations was 68 percent. According to the development profile in 1985, if it were to develop and reach an income level of 1/8, we would expect it to decrease its employment share in manual occupations to 50. That's the expected employment share in manual occupations for a country with that level of income, like Chile. Therefore comparing countries with different levels of income at a point in time involves assessing movements *along* the occupational development profiles.

Despite this simple characterization, these profiles capture most of the variability in occupational employment shares. In 1985, for instance, the R-squared of these three regressions is 0.8 on average, as shown in table 3.1. In 2014, this average decreases to 0.7, mostly driven by the increased dispersion in routine occupational shares.

This analysis, so far, has focused on the employment changes for a given development profile. The next section analyzes the changes of these profiles change over time.

World Polarization

Job polarization, in developed economies, translates to decreases in employment shares for routine occupations. In the sample I use, between 1990 and 2014 the routine employment share for this group of economies (weighted by employment levels) decreased by 0.45 percentage points, annually. Abstract occupations increased its employment share by 0.41, while manual occupations increased it by 0.04.⁶ This is consistent with the evidence in Autor *et al.* (2003), Autor and Dorn (2013), and Goos *et al.* (2014). World polarization, on the other hand, translates to decreases in the development profile for routine occupations, and increases in the development profiles of abstract and manual occupations.

This section focuses on the movements of the occupational development profiles, rather than individual changes in employment shares. Typically, countries with lower income levels have higher manual employment shares, and countries with higher income levels have higher abstract employment shares. The movements along the development profiles suggest different changes in their employment shares as they develop. The countries I analyze are more heterogeneous, and are in different stages of development. Therefore, analyzing the changes in the entire profile provides a way to analyze employment trends, globally.

A second look at Figure 3.1 shows how the world has polarized. For that, we need to compare the development profiles of 1985 (depicted by the black lines) to the profiles of 2014 (depicted by the golden lines). Analogously to the concept of job polarization, world polarization implies that the development profile of routine occupations decreased, while the development profiles of abstract and manual occupations

⁶The list of developed economies is: Austria, Belgium, Denmark, Finland, France, Germany, Greece, Ireland, Italy, Luxembourg, Netherlands, Norway, Portugal, Spain, Sweden, Great Britain, and United States.

increased.

Between 1985 and 2014, the largest change happened in the routine development profile. On average, every decade saw this development profile decrease by 3 percentage points. In 1985, a country with an income of 1/4 of the U.S. 2000 level, like Spain, had in expectation a routine employment share of 38 percent. In 2014, a country with the same income level, like Peru, had in expectation a routine employment share of 27.

This is a novel finding, which contrasts with the analyses in the structural transformation literature. Cross country evidence of employment changes at the *industry* level suggests that currently, a country with an income level of the U.S. in 1980 has a similar industrial composition to what the U.S. had in 1980.⁷ Industries and occupations are closely related, but the connection between development and occupational employment levels changed over time.

Moreover, examining the relationship between polarization and structural transformation more closely reveals that polarization is happening mostly due to changes *within* industries, rather than *between* industries. Intuitively, the routine employment share in a country can decrease because it is moving towards industries that do not require routine workers (structural transformation, or changes between industries), or because all industries are requiring less routine workers (polarization within industries). A shift-share decomposition quantifies the contribution of these two movements. During these years, two thirds of the employment changes happened within industries. This means that polarization is happening in addition to structural transformation, not due to it. Further details of this shift-share analysis are presented in Appendix C.2.

Notice that a lower routine development profile does not mean that all countries

⁷See Herrendorf *et al.* (2014a) for a discussion.

will have lower routine employment shares. It means that for a given level of development, routine employment shares will be lower, but as countries develop, they may have higher routine employment shares. The key distinction are the movements *along* the development profile, and the movements *of* the development profile. Take, for example, Costa Rica's development process. Between 1985 and 2014, its routine employment share barely increased: it went from 29.16 to 29.46 percent.⁸ Its income level, however, did change: in 1985 it was 15 percent of the U.S. level in 2000, and by 2014 it had increased to 31. The development profile of 1985 predicted its routine employment share to increase by 7 percentage points. This corresponds to the movements along the development profile. In 2014, however, the development profiles had changed, and for Costa Rica's income level, it predicted a much lower routine employment share. This corresponds to the movements of the development profile. In practice, these two resulted in a slightly higher routine employment share.

Lastly, notice that the change in the routine development profile is different for countries with low and high income levels. It is not a parallel shift of the profile, and the counterparts of these movements are also different. For countries with lower income levels, this decrease is mostly made up for by a higher manual development profile, and for countries with higher income levels, in abstract. In figure 3.1, the change in the abstract development profile looks more like an increase in its slope, while the change in the manual development profile looks more like an increase in its intercept. This asymmetry implies that, as countries with low income levels develop, manual occupation shares will decrease, but more slowly than the development profile suggests.

To summarize the empirical findings, in this section I introduce the concept of the

⁸During this period, most of the labor reallocation happened between abstract and manual occupations, exchanging about 6 percentage points.

occupational development profiles. Its goal is to link the distribution of occupational employment shares to income levels. These development profiles have shifted over time, resulting in world polarization: a lower development profile for routine occupations, and higher profiles for manual and abstract occupations. This is, modern growth has been biased against routine occupations. The following section quantifies task-specific productivity levels that are compatible with these facts, and what their changes imply for the following years.

3.3 Further Polarization Is on Its Way

The way an economy allocates its resources is informative of its productive technology. In this section, I apply this idea to the occupational distributions through a development accounting exercise. This allows to quantify the task-specific productivities behind occupational distributions, and document the biases in technological progress. This section is structured as follows. The first part introduces the polarization accounting framework, and the way to map it to the data. The second explains its results; how technical change has been biased taking as a benchmark the 1985 development profiles. The third part explains the VAR framework I use to forecast technological trends, and the fourth section analyzes these forecasts until 2050.

3.3.1 *Polarization Accounting*

The goal of this section is to provide a way to quantify two aspects of modern economic growth: the technological factors behind occupational employment shares, and their change over time. This is, I perform a polarization accounting exercise, similar to the development accounting discussed in Hsieh and Klenow (2010). Development accounting “uses cross-country data on output and inputs, at one point in time, to assess the contribution of differences in factor quantities and the efficiency

with which these factors are used” (Caselli, 2005, p. 681). Polarization accounting uses data on occupational employment shares to quantify these efficiency levels.

For that purpose, consider an economy following the task-based approach to production, as in Acemoglu and Autor (2011). To produce output, firms need to combine tasks according to a production function. In that sense, tasks are the basic building blocks in production. To map this framework to the data discussed in the previous section, I assume there are three tasks: abstract (a), routine (r), and manual (m). For notation purposes, task-specific variables are denoted by the j subindex, so that $j \in \{a, r, m\}$.

The production of task j depends on two components: an amount labor, $l_{c,j,t}$, and a task-specific productivity level $A_{c,j,t}$. Their product results in efficiency units of labor. The subindex c represents the country, and t the time. Technology is labor-augmenting, and these task-specific productivity levels effectively capture all of the factors increasing these efficiency units. This is a reduced-form way of grouping factors like equipment and machines that are specific to the production of tasks, and capital stocks and total factor productivity that are not.⁹

Final output requires these three tasks. Production happens according to the following constant elasticity of substitution production function:

$$y_{c,t} = \left[\sum_{j \in \{a,r,m\}} \omega_j^{\frac{1}{\varepsilon}} (A_{c,j,t} l_{c,j,t})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \quad (3.2)$$

where $\varepsilon > 0$ is the elasticity of substitution among tasks, and $\omega_j > 0$ is the production intensity of task j . These task intensities add up to one.¹⁰ In this framework, task-

⁹An alternative, but equivalent approach would look at the costs of producing tasks, rather than its efficiency levels like Goos *et al.* (2014).

¹⁰These intensities are raised to the power $1/\varepsilon$ so that the limit case where $\varepsilon \rightarrow 0$ converges to a Leontief utility function: $\lim_{\varepsilon \rightarrow 0} = \min_{j \in \{a,r,m\}} \{\omega_j A_{c,j,t} l_{c,j,t}\}$.

specific productivity levels and the distribution of labor are country-specific; the elasticity of substitution and the task intensities are not.

Labor can be allocated to the production of any of these tasks. For simplicity (and data limitations), labor is homogeneous and has no task specificity to it. Since the main interest is on the technological changes leading to polarization, I assume that labor is perfectly mobile across tasks.¹¹ I normalize the total labor force to 1, so that these labor inputs represent occupational employment shares. This means that $y_{c,t}$ effectively corresponds to output *per worker*. This is also the measure I match in the data. Then,

$$1 = \sum_{j \in \{a,r,m\}} l_{c,j,t} \quad (3.3)$$

I model the allocation of labor through competitive markets. Firms are price and wage takers, and there is free entry into the production of the final good. Their production technology is represented the production function (C.12), so that their optimization problem is:

$$\max_{\{l_{c,j,t}\}_{j \in \{a,r,m\}}} \left[\sum_{j \in \{a,r,m\}} \omega_j^{\frac{1}{\varepsilon}} (A_{c,j,t} l_{c,j,t})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} - \sum_{j \in \{a,r,m\}} w_{c,j,t} l_{c,j,t} \quad (3.4)$$

In this setting, the final good is the numeraire, and firms hiring labor face wages $w_{c,j,t}$. Free entry in this context implies wage equalization, since labor is homogeneous and mobility is costless. In addition, profits are zero due to constant returns to scale

¹¹This is a common assumption in the structural transformation literature. See, for example, Baumol (1967), Kongsamut *et al.* (2001a), Ngai and Pissarides (2007b), and Duarte and Restuccia (2010). Their focus, as in this paper, is on the technological aspects of employment movements, so this assumption is mostly made out of convenience. Other papers, like Caselli and Coleman (2001a) and Bárány and Siegel (2018) consider labor heterogeneity. In a separate project, we analyze the costs of reallocation.

in production. The optimal allocation of labor depends on the production intensities, productivity levels, and elasticity of substitution as follows:

$$\frac{l_{c,j,t}}{l_{c,k,t}} = \frac{\omega_j}{\omega_k} \left(\frac{A_{c,k,t}}{A_{c,j,t}} \right)^{1-\varepsilon} \quad (3.5)$$

With this optimality condition and the normalization of the labor force, we can aggregate the production function. Final output can be expressed exclusively as function of the productivities in each of these tasks:

$$y_{c,t} = \left[\sum_{j \in \{a,r,m\}} \frac{\omega_j}{A_{c,j,t}^{1-\varepsilon}} \right]^{\frac{1}{\varepsilon-1}} \quad (3.6)$$

These last two expressions are key to matching the model with the data. For country c , we observe its income level $y_{c,t}$, and its distribution of employment by occupations, $\{l_{c,j,t}\}_{j \in \{a,r,m\}}$. The objective of this framework is to infer the task-specific productivities, $A_{c,j,t}$. Combining equations (3.5) and (3.6) allows to do so by setting a system of 3 equations in 3 unknowns:

$$A_{c,j,t} = y_{c,t} \left[\sum_{k \in \{a,r,m\}} \omega_k \frac{l_{c,k,t}}{l_{c,j,t}} \right]^{\frac{1}{1-\varepsilon}} \quad (3.7)$$

I set the elasticity of substitution ε to 0.35, as in Vindas (2017), which analyzes occupational U.S. data. In addition, I normalize income levels and task-specific productivities to U.S. levels in 2000, which provides a way to estimate the task-intensity parameters ω_j . Therefore, both income and task-specific productivities are expressed in relative terms to U.S. levels in 2000, maintaining consistency with the units of measurement of the empirical section.¹²

The last part of this article forecasts task-specific productivity growth as a way to forecast the development profiles. For that, we need an expression for the occupational

¹²This normalization plays no important role in the following analyses since it scales productivities by the same factor.

employment shares as a function of the task-specific productivities. This results from combining equations (3.3) and (3.5):

$$l_{c,j,t} = \frac{\omega_j / A_{c,j,t}^{1-\varepsilon}}{\sum_{k \in \{a,r,m\}} \omega_k / A_{c,k,t}^{1-\varepsilon}} \quad (3.8)$$

3.3.2 Accounting Results

Countries develop and change their occupational employment shares according to their task-specific productivities. This section documents the productivity changes inferred from the polarization accounting exercise. The first part documents the historical changes by showing their growth rates. The second compares them to growth patterns implied by the development profiles of 1985.

Historical Growth

Table 3.2 shows the averages and the standard deviations of the productivity growth rates between 1985 and 2014. Each country-year pair is counted independently, so that countries with longer histories have a larger weight. It also presents these summary statistics, but weighted by their employment levels.

Between 1980 and 2014, productivity growth rates were the lowest in abstract occupations, independently of whether these are weighted by employment or not. This is a clear reflection of Baumol's cost disease, but applied to occupations (Baumol *et al.*, 1985). Routine and manual occupations have very similar productivity growth rates. Productivity in routine tasks is slightly higher at the country level, but weighting by employment reverts this.

These changes in task-specific productivities relate to occupational employment shares through equation (3.5):

$$\frac{l_{c,j,t}}{l_{c,k,t}} = \frac{\omega_j}{\omega_k} \left(\frac{A_{c,k,t}}{A_{c,j,t}} \right)^{1-\varepsilon}$$

Table 3.2: World Task-Specific Productivity Growth Rates (%): 1980-2014

	N	Unweighted		Employment-weighted	
		Mean	Std. Dev.	Mean	Std. Dev.
Abstract	2186	1.027	7.732	1.885	5.287
Routine	2186	4.029	6.741	3.917	4.765
Manual	2186	3.840	6.901	3.933	4.573

Task-specific productivity growth rates estimated from equation (3.7), not average log-differences. Unweighted statistics consider each country-change pair equally; employment-weighted consider the country's employment levels

Source: author's calculations using IPUMS, ILOSTAT & PWT data.

Over time, employment shares will depend on the relative task-specific productivities, and their elasticity of substitution. In the production function, the magnitude of ε determines how employment shares respond to changes in relative productivities. When $\varepsilon > 1$, tasks in production are good substitutes: if $A_{c,j,t}$ increases, its employment share will do so as well. When $\varepsilon = 1$, the production function converges to a Cobb-Douglas, and employment shares do not depend on the productivities. When $\varepsilon < 1$, as I assume here, tasks are complements, and the production function resembles more a Leontief technology. When $A_{c,j,t}$ increases, this will actually decrease the employment share in occupation j .

These summary statistics are a good starting point to analyze how productivities changed over time. The nonlinearities in the production function mean that these effects will be different, depending on the productivity levels of the other tasks. The following section, then, analyzes productivity growth patterns taking into account the development profiles.

Measuring Productivity Biases

The development profiles summarize the link between levels of income and occupational employment shares in a given year. From equation (3.1):

$$d_{j,t}(y) = \beta_{0,j,t} + \beta_{1,j,t}y + \beta_{2,j,t}y^2$$

Therefore, these profiles summarize the expected path of occupational employment shares as countries develop. These also have a counterpart for the path of task-specific productivities. From equation (3.7), denote these expected productivity paths by \mathcal{A}_j :

$$\mathcal{A}_j(y; t) = y \left[\sum_{k \in \{a,r,m\}} \omega_j \frac{d_{k,t}(y)}{d_{j,t}(y)} \right]^{\frac{1}{1-\varepsilon}}$$

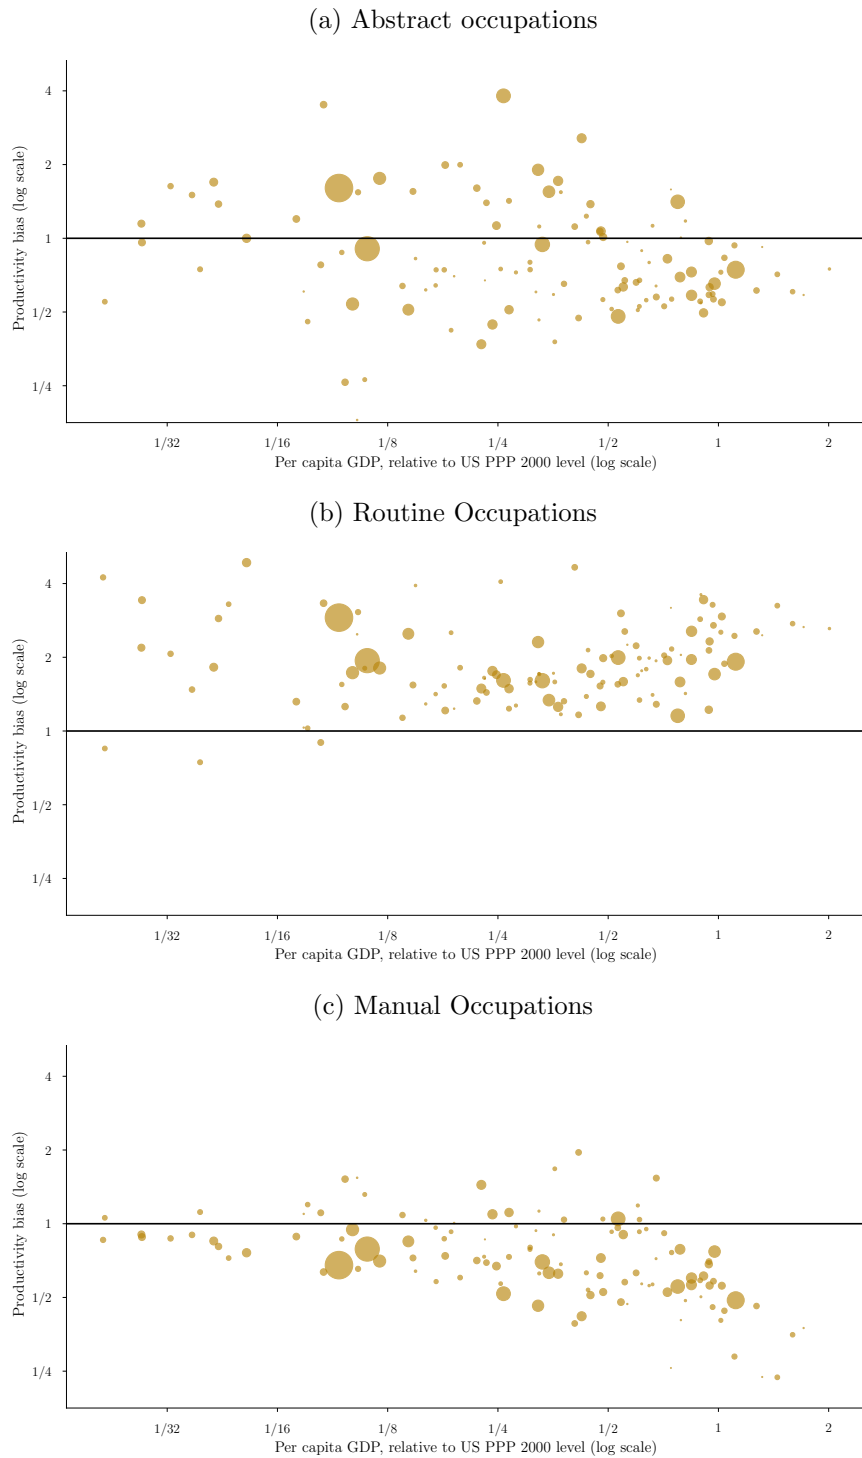
Mechanically speaking, for all levels of income y , the development profiles $d_{j,t}(y)$ predict a distribution of employment shares, which altogether imply a level of task-specific productivity. The last section documented how these development profiles have shifted over time. Then, the changes in these paths of productivities provide a natural benchmark to evaluate whether technological growth has been biased.

Define the productivity bias in task j with respect to period t as the log-difference in a measured task-specific productivity level, A_j , with respect to the productivity level that the development profile of period t predicts for its income level, $\mathcal{A}_j(y; t)$. Denote this productivity bias by $b_j(y, A_j; t)$. Then,

$$b_j(y, A_j; t) = \log(A_j) - \log(\mathcal{A}_j(y; t)) \tag{3.9}$$

In this definition, productivity is biased against occupations j if $b_j(y, A_j; t) > 0$. This is, if task-specific productivity is higher than the development profile in period t predicted for the income level y . The bias is *against* occupation j because a higher task-specific productivity level implies a *lower* employment share. Conversely, productivity is biased in favor of occupation j if $b_j(y, A_j; t) < 0$ by the same reasoning.

Figure 3.2: Productivity Biases in 2014 with respect to 1985



Source: author's calculations using ILOSTAT, IPUMS & PWT.

Figure 3.2 shows the productivity biases in 2014 with respect to 1985. The horizontal axes show income levels, and the vertical axes the biases measured through (3.9). As before, each circle represents a country, and its size is proportional to its employment level.

Between 1985 and 2014, technical change has consistently been biased against routine occupations, and in favor of manual occupations. The bias in abstract occupations, shown in the first panel, follows no consistent pattern: for some countries it has been positive, and for some it has been negative. There does not seem to be any systematic relationship with respect to income levels, either. For routine occupations, the result is markedly different: productivity growth has been biased against these occupations because it has consistently been positive. Only in three countries the bias has been negative. Manual occupations show the opposite pattern: their productivity growth patterns result in negative biases, producing higher manual employment shares than the development profile of 1985 suggests. Furthermore, this bias shows a negative relationship with income, so that countries with higher levels of income have shown larger biases in favor of manual occupations. Then, technical change has consistently been biased against routine occupations, and in favor of manual occupations.

3.3.3 VAR Analysis

The previous sections focused on a historical analysis of task-specific productivities, by quantifying its growth patterns and biases. The following sections provide a forward-looking exercise by extrapolating these productivity trends, and analyzing the implications for employment distributions in the future.

A forecast of the world distribution of occupations could be based on the latest development profile, and predict employment changes *along* these profiles. This

wouldn't be a satisfactory approach: the productivity biases between 1985 and 2014 suggest that movements *of* the development profiles are important in describing how the world distribution of occupations has changed. This means that an exercise in forecasting should focus on the evolution of the productivities, and their continued departure from the established development profiles. In this section, I analyze task-specific productivity growth through a VAR model. This is a flexible enough framework allowing for different growth rates across countries and task-specific productivities, as well as interaction terms.

The type of analysis that I use here has commonly been applied in studies of cross-country convergence of per capita GDP levels, like Barro and Sala-i-Martin (1992) and Caselli *et al.* (1996). This line of thought considers a negative relationship between initial income levels and their subsequent growth path. My analysis differs in two dimensions. First, instead of considering one level of income, I study a three-dimensional vector of productivities. These productivities relate to income levels through equation (3.6). Second, the productivities I study are expressed as *gaps* from a technological frontier. This follows the idea of technological diffusion in Parente and Prescott (1994).

In this particular setting, I use the U.S. productivities $A_{US,j,t}$ as the technological frontiers. I forecast future task-specific productivity levels in two steps. In the first, I extrapolate historical trends for the U.S. productivity levels. In the second, I use a VAR model to predict the path of the gaps for each country relative to the U.S. These gaps are defined as:

$$\tilde{A}_{c,j,t} = \frac{A_{c,j,t}}{A_{US,j,t}} \tag{3.10}$$

The growth process of the U.S. productivities follows a simple path of constant

Table 3.3: U.S. Productivity Log-Differences: 1980-2014

	Occupational Productivity Log-Difference		
	Abstract	Routine	Manual
Constant	0.0011 (0.005)	0.0488*** (0.005)	0.0236*** (0.007)

OLS estimates of equation (3.11). Asterisks indicate statistical significance at 10 percent (*), 5 percent (**), and 1 percent (***).

Source: author's calculations using IPUMS, ILOSTAT & PWT data.

growth rates. The model I estimate is:

$$\Delta \log(A_{US,j,t}) = \alpha_{US,j} + \epsilon_{US,j,t} \quad (3.11)$$

Which is a trio of random walks with drifts. Table 3.3 shows the results of regressing the difference of the U.S. log-productivity levels with only a constant term.

The average difference in the log-productivity of routine tasks is the highest at 0.049 log points per year, followed by manual at 0.024. These two are statistically significant. Productivity changes in abstract occupations are much lower at 0.001 log points per year. This is not statistically different from zero; its standard error is quantitatively similar to the other two, but its level is too low. These are the differences I use to forecast the technological frontier up to 2050. Between 2014 and 2050, productivity in abstract tasks will grow by a factor of 1.03, while in routine and manual the growth factor is 3.38 and 1.80, respectively.

The second step involves a VAR forecast of the productivity *gaps* with respect to the U.S. levels. Because I study the three productivities jointly, I group them into a

vector of log-gaps: $\tilde{\mathbf{a}}_{c,t} = [\log \tilde{A}_{a,c,t}, \log \tilde{A}_{r,c,t}, \log \tilde{A}_{m,c,t}]'$. The model I estimate is:

$$\Delta \tilde{\mathbf{a}}_{c,t} = \boldsymbol{\alpha}_c + \mathbf{B} \tilde{\mathbf{a}}_{c,t-1} + \mathbf{e}_{c,t} \quad (3.12)$$

The dependent variables are the differences in the productivity log-gaps, $\Delta \tilde{\mathbf{a}}_{c,t}$. Countries can vary in terms of their institutions and resources, which can result in different long-run productivity gaps. I capture these differences through country fixed effects, grouped in the vector $\boldsymbol{\alpha}_c = [\alpha_{a,c}, \alpha_{r,c}, \alpha_{m,c}]$. The 3×3 matrix \mathbf{B} contains the coefficients associated with the lagged log-gaps, and $\mathbf{e}_{c,t} = [\varepsilon_{a,c,t}, \varepsilon_{r,c,t}, \varepsilon_{m,c,t}]$ is the vector of error terms that are centered around zero, and are identically and independently distributed.

Two properties of this system are of interest: the long-run expected productivity log-gaps, and the dynamic path towards it. The long-run expected productivity log-gaps are the equivalent to a steady-state gap. These follow from equation (3.12), and require the expected differences in log-productivities to be zero. Therefore:

$$\tilde{\mathbf{a}}_{c,t} = \tilde{\mathbf{a}}_{c,t-1} = \bar{\mathbf{a}}_c \quad (3.13)$$

$\bar{\mathbf{a}}_c$ is the vector of long-run expected productivity log-gaps. This is equal to:

$$\bar{\mathbf{a}}_c = -\mathbf{B}^{-1} \boldsymbol{\alpha}_c \quad (3.14)$$

Notice these gaps have a common component through \mathbf{B} , but also depend on the country-specific fixed effects.

We now turn to the dynamic path towards these long-run log-gaps. For that, it is useful to rewrite equation (3.12) as deviations from the long-run productivity log-gaps. In expectation:

$$(\tilde{\mathbf{a}}_{c,t} - \bar{\mathbf{a}}_c) = (\mathbf{I} + \mathbf{B})(\tilde{\mathbf{a}}_{c,t-1} - \bar{\mathbf{a}}_c) \quad (3.15)$$

Then, the deviations from the long-run productivity log-gaps are fully determined by $\mathbf{I}+\mathbf{B}$. In particular, this implies that the forecast k periods in advance is:

$$(\tilde{\mathbf{a}}_{c,t+k} - \bar{\mathbf{a}}_c) = (\mathbf{I}+\mathbf{B})^k(\tilde{\mathbf{a}}_{c,t} - \bar{\mathbf{a}}_c) \quad (3.16)$$

This is the expression that I use to forecast the productivity log-gaps. Whether this system is convergent, and its speed of convergence depends on the eigenvalues of $\mathbf{I}+\mathbf{B}$.

Table 3.4 presents three components of the estimated VAR system. First, the estimates of the \mathbf{B} matrix, second the eigenvalues of the matrix $\mathbf{I}+\mathbf{B}$, which dictate the dynamics over time, and lastly, the rates of convergence to the long-run expected productivity log-gaps.

Most of the estimates in \mathbf{B} are statistically significant. The columns contain the estimates for each of the dependent variables, the differences of the log-productivity gaps. The rows group the effects of the independent variables, the lagged productivity log-gaps. The diagonal terms in abstract and routine productivities are negative, as expected. Heuristically, if we ignore the non-diagonal terms, countries with higher deviations from their long-run productivity log-gaps in abstract and routine tasks close these deviations faster. This is not the case for manual productivity, since its difference is associated with a positive, but not statistically significant, estimate of its lagged log-gap. This means that, log-gaps in manual productivity are mostly driven by log-gaps in routine productivity since it's the only coefficient that is statistically significant.

The dynamics of a growth regression are simple to interpret if the dependent variable is unidimensional. The model can be rearranged as a difference equation, and the coefficient associated to the first lag determines the speed of convergence to

Table 3.4: VAR Estimate Results

Productivity log-gap in	Log-difference of productivity gap in		
	Abstract	Routine	Manual
Abstract	-0.087*** (0.013)	0.015 (0.010)	-0.001 (0.010)
Routine	-0.034*** (0.009)	-0.100*** (0.010)	-0.061*** (0.009)
Manual	0.051*** (0.009)	0.050*** (0.007)	0.013 (0.008)
Eigenvalues associated to the dynamic system			
	0.915	0.955+0.019i	0.955-0.019i
Annual convergence rates (moduli)			
	0.085	0.045	

Upper section contains the VAR estimate of matrix \mathbf{B} in (3.12). Middle section contains the eigenvalues of matrix $\mathbf{I}+\mathbf{B}$, and lower section contains their moduli. Asterisks indicate statistical significance at 10 percent (*), 5 percent (**), and 1 percent (***).

Source: author's calculations using IPUMS, ILOSTAT & PWT data.

its steady state.¹³ A three-dimensional vector requires a slightly different approach. The system (3.15) has to be diagonalized to get an autonomous system of difference equations. This simplifies greatly the analysis of its dynamic properties, and the steps to do so are explained in appendix C.3. This analysis boils down to two sets of values: the eigenvalues of the system, and their moduli. These are presented in the bottom part of table 3.4.

The system, altogether, is convergent with small cycles. The cyclical behavior follows from the eigenvalues, and the convergence from the moduli of these eigenvalues. Complex eigenvalues, like the ones from this system, indicate cyclical behavior. Convergence to the long-run log-gaps is not monotone, but the cyclical components (i.e., the coefficients associated to the imaginary part) are fairly small. The moduli of these eigenvalues are less than one, which mean that the system is convergent to the long-run productivity log-gaps in (3.14). Notice that these convergence rates correspond to the diagonalized system, which is a linear combination of the three productivity log-gaps. These converge at a rate between 4.5 and 8.5 percent per year. The non-diagonalized convergence rates, the ones at the log-gap level, are discussed in the following section. These are lower, and closer to the estimates of Sala-i-Martin (1994).

3.3.4 *Forecasting Results*

The previous section analyzed the growth patterns of task-specific productivities through a VAR model. In this one, I use the estimated model to extrapolate them. This exercise forecasts historical trends for 36 years, and the main goal is to describe the occupational development profiles in 2050. The forecasts begin in 2015, since current income data in the Penn World Tables is published up to 2014.

¹³See Barro and Sala-i-Martin (2004) for several applications of this analysis.

Table 3.5: U.S. Task-Specific Log-Productivity & Occupational Employment Shares (1980, 2014 & 2050)

Year	Log-Productivity			Employment Shares		
	Abstract	Routine	Manual	Abstract	Routine	Manual
1980	15.305	14.509	14.662	29.42	45.94	24.64
2014	15.341	16.167	15.465	42.66	32.06	25.29
2050	15.379	17.926	16.316	56.90	19.71	23.39

Historical data for 1980 & 2014; 2050 forecasts from model (3.11).

Source: author's calculations using IPUMS, ILOSTAT & PWT data.

Even though the main goal is to describe employment shares, this exercise forecasts task-specific productivities. Employment shares, as described in equation (3.8) are non-linear functions of the countries' task-specific productivities. Because of that, I focus on predicting these productivities as a way to construct occupational employment shares. This, as explained previously, involves two steps. In the first one, I forecast task-specific productivity growth in the U.S., which is the technological frontier in this framework. The second one forecasts the gaps with respect to that frontier. Ultimately, I combine these two forecasts to get employment shares, and describe the occupational development profiles.

Technological Frontier Forecast

The results of the productivity extrapolation in the U.S. are in Table 3.5. It contains the task-specific log-productivities and their implied employment shares in 1980, 2014, and 2050. Naturally, we can compare two periods. The first contains historical U.S. data from 1980 to 2014, and the second contains the forecasted data from 2014 to 2050.

Figure 3.3: U.S. Occupational Employment Shares: 1980 - 2050



Source: author's calculations using ILOSTAT, IPUMS & PWT.

By construction, log-productivities in the forecasted years change at the same rate as the historical data. The occupational employment shares, by contrast, change differently. The employment share in routine occupations continues to decrease, but its change slows down. During the historical period it decreased by 0.41 percentage points, annually, and in the forecasted it decreases by 0.34. The employment share in manual occupations goes from minor increases to minor decreases. During the historical period it increased by 0.02 percentage points annually, and decreases by 0.05 in the forecasted data. The increases in the employment share in abstract occupations slightly accelerate from 0.39 percentage points, to 0.40 in the forecasted data. These differences are due to the non-linearities that determine the employment shares, even though the underlying productivities grow at constant rates. These shares are also plotted in figure 3.3. By the end of the forecasted period, routine occupations will have the lowest employment share. Manual occupations will decrease during these years, but routine occupations will decrease at a much faster pace.

Productivity Gaps Forecast

The second step forecasts the differences in log-productivity gaps with the VAR model. As discussed in the previous section, the dynamics imply a convergent path towards long-run productivity log-gaps, which are specific to each country in the sample. Table 3.6 shows the results of this forecast. To summarize the information at the world level, it reports averages across countries. First, it contains the productivity log-gaps with respect to the technological frontier in 2014 and 2050. These follow from the definition in equation (3.10): $\tilde{\mathbf{a}}_{c,t} = [\log \tilde{A}_{a,c,t}, \log \tilde{A}_{r,c,t}, \log \tilde{A}_{m,c,t}]'$. It also shows the log-deviation from the long run productivity gaps: $\tilde{\mathbf{a}}_{c,t} - \bar{\mathbf{a}}_c$. The last column shows the average change over this period, which is the same for the log-gaps and the log-deviations. Log-gaps show how far countries are from the technological frontier; log-deviations show how far countries are from their long-run productivity gaps.

The larger log-gaps are in the productivity of manual tasks, followed by routine and abstract tasks. These log-gaps close over time, and maintain the same order by the end of the forecasted period. This is true as well for the log-deviations from the long-run gaps. Notice, however, that these log-deviations are not negligible by the end of the forecasted period. For abstract occupations, a log-deviation of -0.142 translates to a ratio of its productivity gap with respect to its long-run productivity gap of 0.9.¹⁴ This is, by 2050, the average productivity gap in abstract tasks will be at 90 percent of its long-run value. For routine and manual tasks, these values will be 81.3 and 80.7 percent. This means that the deviations from the long-run productivity gaps will still remain an important determinant of the employment distributions. The last column shows the average change during the forecasted period, which is the same

¹⁴These are base 2 logarithms, so that elevating 2 to the power \tilde{a} results in the ratio of the productivity levels.

Table 3.6: World Productivity Gaps: 2014 & 2050

Occupation	Log-gaps		Log-deviations		Average
	2014	2050	2014	2050	change
Abstract	-0.503	-0.013	-0.637	-0.142	0.014
Routine	-1.231	-0.519	-1.016	-0.299	0.020
Manual	-3.007	-1.961	-1.363	-0.309	0.029

Log-gaps as defined in equation (3.10), and log-deviations as in (3.15).

Averages correspond to the simple mean across countries; average change is the same for log-gaps and log-deviations.

Source: author's calculations using IPUMS, ILOSTAT & PWT data.

for the log-gaps and the log-deviations. Overall, convergence to the long-run log-gaps is fastest in manual occupations, followed by routine and abstract.

Employment Shares Forecast

The forecast of the productivity log-gaps and the productivity frontier provides enough information to forecast the employment distributions up to 2050. For that, I use equation (3.8):

$$l_{c,j,t} = \frac{\omega_j / A_{c,j,t}^{1-\varepsilon}}{\sum_{k \in \{a,r,m\}} \omega_k / A_{c,k,t}^{1-\varepsilon}}$$

Table 3.7 shows yearly average changes in occupational employment shares for two periods: the historical period between 1980 and 2014, and the forecasted period between 2014 and 2050. To summarize the changes in the world distribution of occupations, it shows two averages: one that weights equally each observation, and another one that weights them according to the countries' employment levels in 2014.

The occupational changes observed during the historical period will continue until 2050. Abstract occupations will continue to increase, while routine and manual will

Table 3.7: World Annual Occupational Employment Share Changes: 1980-2014 & 2014-2050

Occupation	Country average (unweighted)		World average (weighted)	
	1980-2014	2014-2050	1980-2014	2014-2050
Abstract	0.346	0.397	0.195	0.294
Routine	-0.193	-0.243	-0.082	-0.211
Manual	-0.154	-0.154	-0.113	-0.083


Changes expressed as percentage points per year. Country averages weight each country-year change equally; weighted averages use employment in 2014.

Source: author's calculations using IPUMS, ILOSTAT & PWT data.

continue to decrease. Whether we look at the simple averages, or the employment-weighted ones, the results are very similar. Worldwide, the employment-weighted share in abstract occupations will increase, and will accelerate over time. The historical data shows that it increased by 0.195 percentage points every year, and the forecasted shares increase by 0.294 percentage points per year. The worldwide decrease in routine occupations will accelerate quite dramatically. From decreasing 0.082 percentage points annually, the forecast establishes that it will decrease 0.211 in the following years. Finally, the drop in manual occupations will slow down. Historically, it fell by 0.113 percentage points annually, but the forecasts suggest it will drop by 0.083 annually. These changes may seem small at face value, but represent enormous amounts of workers reallocating. Over the course of the forecasted period, 226 million workers will reallocate to abstract occupations, 162 millions will move out of routine occupations, and 64 millions will move out of manual occupations.

Graphically, these changes are plotted in figure 3.4. It shows the occupational employment shares, weighted by the countries' employment levels, between 1985 and

Figure 3.4: World Occupational Employment Shares: 1985 - 2050



Source: author’s calculations using ILOSTAT, IPUMS & PWT.

2050.¹⁵ At the world level, abstract occupations increase and manual occupations decrease monotonically, but routine occupations show a “hump-shape”. This is reminiscent of the analyses of structural transformation (employment shares by industry, rather than occupation), like Ngai and Pissarides (2007b), Buera and Kaboski (2009), and Duarte and Restuccia (2010). Differently to these, this plot aggregates the employment shares at the world level, instead of the country level data. Qualitatively, it shares the rise and fall of one of the categories, in this case routine occupations. The world employment share in abstract occupations will increase from 20 percent in 2014 to 31 in 2050. Routine occupations will decrease from 23 percent to 15, and manual occupations will decrease slightly from 57 to 54 percent.

The last paragraphs described the predicted changes at the world level, aggregating over countries. These countries, however, will still differ in their levels of income and employment distributions. Precisely due to these differences, this paper analyzed

¹⁵This plot begins in 1985 since between 1980 and 1984, employment for the countries in the sample reached less than 65% of the level of the full sample.

the occupational development profiles, rather than world totals. Figure 3.5, then, shows the predicted development profile of 2050, in addition to the development profiles of 1985 and 2014. As before, the observations in black represent data from 1985, and the observations in gold represent data from 2014. The forecasted data for 2050 is presented in maroon.

The main result is that world polarization will continue. Between 1985 and 2014, this meant a lower development profile for routine occupations, and higher development profiles for abstract and manual occupations. Between 2014 and 2050, this will be the case as well: the development profile in abstract occupations will be slightly higher, in routine occupations it will be much lower, and in manual occupations it will be higher. Most of the changes will happen among these last two occupations.

Perhaps surprisingly, the development profile of abstract occupations changes little, compared to the change between 1985 and 2014. The shape of the profile in 2050, in the first panel, shows only a slight increase. This is a reflection of the biases documented in figure 3.2. In abstract tasks, productivity growth showed no systematic bias. Because of that, the development profile changed little.

The development profile in routine occupations, in the second panel, keeps decreasing. It also shows a larger drop than between 1985 and 2014. This means that in the following years, productivity growth in routine tasks will be high enough to make many routine workers redundant. It's also worthy to note how the "hump-shape" flattens over time. During the forecasted period, countries will catch up to the technological frontier, that is making routine tasks very productive. This causes technological progress in the rest of the countries biased against routine occupations, resulting in a flatter profile.


For manual occupations, in the third panel, changes happen in the opposite direction. The development profile shifts upwards, which means that for a given level of

Figure 3.5: Occupational Development Profiles: 1985, 2014 & 2050


(a) Abstract occupations



(b) Routine Occupations



(c) Manual Occupations



Source: author's calculations using ILOSTAT, IPUMS & PWT.

income, a country will have on average a higher employment share in 2050 compared to both 1985 and 2014. The reasoning behind this change in the development profile, is the same as for routine occupations. The productivity growth bias, however, happens in the opposite direction.

These results are driven completely by the growth patterns of task-specific productivities. These, as explained before, are built in two steps: one forecasts the evolution of the technological frontier, and the other one forecasts the gaps with respect to that frontier. Which one is driving most of the employment shifts? Over time, we can decompose overall productivity changes into these two components:

$$\Delta \mathbf{a}_{c,t} = \Delta \mathbf{a}_{US,t} + \Delta \tilde{\mathbf{a}}_{c,t} \quad (3.17)$$

This tells us how much of the changes in productivity is due to the frontier growing, and how much by the gaps closing. Table 3.8 shows this decomposition. The first column shows the average annual change in the task-specific productivities, $\Delta \mathbf{a}_c$, and breaks it down into the contribution of the technological frontier, $\Delta \mathbf{a}_{U.S.}$, and the change in countries' gaps, $\Delta \tilde{\mathbf{a}}_c$.

The contribution of these two factors, overall, is sizable. This decomposition, however, throws very different results across tasks. For abstract tasks, most of the growth is due to the gaps closing, while in routine tasks, the technological frontier growth outpaces the contribution of the gaps closing. For manual tasks, the growth in the productivity still plays a predominant role. This means that the obsolescence of routine workers is mostly a product of the productivity changes happening at the technological frontier. Productivity growth is so high that its effect will eventually be transmitted to the rest of the countries, and push routine employment shares down, as in this forecast exercise.

Table 3.8: World Productivity Growth Sources

	Average log-productivity	Average contribution (%)	
	change (annual)	Frontier	Gaps
Abstract	0.013	17.2	82.8
Routine	0.067	74.2	25.8
Manual	0.050	55.2	44.8

Results of decomposition (3.17). Averages correspond to unweighted means across countries.

Source: author’s calculations using IPUMS, ILOSTAT & PWT data.

3.4 Conclusions

In this paper, I present new stylized facts about the distribution of occupations in the world. I expand significantly the countries in the analysis to a sample of 119 countries covering all levels of economic development. The result of this analysis is that job polarization is a global phenomenon.

At any point in time, there is a strong link between a country’s development level and its occupational employment shares. This is what I call the *occupational development profile*. Over time, this profile has shifted, resulting in world polarization. The development profile for routine occupations has decreased, coupled with higher manual and abstract development profiles. The modern growth experience, then, is biased against routine occupations.

To analyze these development patterns, I follow the grouping principle for occupations and develop a polarization accounting framework at the *task* level. Technical change has been biased against routine occupations and biased in favor of manual occupations.

These technological trends imply that as countries continue to develop, we can

expect lower employment shares in routine occupations, and higher in manual occupations, worldwide. Through a vector autoregression (VAR) analysis, I forecast the path of productivity growth, and conclude that world polarization will continue. In the following years, then, the development profile for routine occupations will keep decreasing, and the profile for manual occupations will keep increasing.

In this paper, I present a first analysis of the task-specific productivities behind world polarization. Future research plans include the expansion of this model to endogenize technical progress through an *task-investment-specific* growth model. The details of such a framework are presented in Appendix C.4, and its implementation is left as future work.

REFERENCES

- Acemoglu, D. and D. Autor, *Skills, Tasks and Technologies: Implications for Employment and Earnings*, vol. 4 of *Handbook of Labor Economics*, chap. 12, pp. 1043–1171 (Elsevier, 2011), URL <https://ideas.repec.org/h/eee/labchp/5-12.html>.
- Acemoglu, D. and V. Guerrieri, “Capital deepening and nonbalanced economic growth”, *Journal of Political Economy* **116**, 3, 467–498 (2008).
- Aguiar, M. and E. Hurst, “Measuring trends in leisure: The allocation of time over five decades”, *The Quarterly Journal of Economics* **122**, 3, 969–1006, URL <http://www.jstor.org/stable/25098866> (2007).
- Autor, D., “The Polarization of Job Opportunities in the U.S. Labor Market: Implications for Employment and Earnings”, Tech. rep., The Hamilton Project and Center for American Progress (2010).
- Autor, D. and D. Dorn, “The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market”, *American Economic Review* **103**, 5, 1553–97, URL <https://ideas.repec.org/a/aea/aecrev/v103y2013i5p1553-97.html> (2013).
- Autor, D., F. Levy and R. Murnane, “The Skill Content of Recent Technological Change: An Empirical Exploration”, *The Quarterly Journal of Economics* **118**, 4, 1279–1333, URL <https://ideas.repec.org/a/oup/qjecon/v118y2003i4p1279-1333.html> (2003).
- Autor, D. H., D. Dorn and G. H. Hanson, “The china syndrome: Local labor market effects of import competition in the united states”, *American Economic Review* **103**, 6, 2121–68, URL <http://www.aeaweb.org/articles?id=10.1257/aer.103.6.2121> (2013).
- Bárány, Z. L. and C. Siegel, “Job polarization and structural change”, *American Economic Journal: Macroeconomics* (2018).
- Barro, R. and X. Sala-i-Martin, *Economic Growth* (MIT Press, 2004), second edn.
- Barro, R. J. and X. Sala-i-Martin, “Convergence”, *Journal of Political Economy* **100**, 2, 223–251, URL <http://www.jstor.org/stable/2138606> (1992).
- Baumol, W., “Macroeconomics of unbalanced growth: The anatomy of urban crisis”, *The American Economic Review* **57**, 3, 415–426 (1967).
- Baumol, W., S. A. B. Blackman and E. N. Wolff, “Unbalanced growth revisited: Asymptotic stagnancy and new evidence”, *The American Economic Review* **75**, 4, 806–817, URL <http://www.jstor.org/stable/1821357> (1985).
- Baumol, W. J. and W. G. Bowen, *Performing Arts - The Economic Dilemma* (MIT Press: New York, NY, 1968).

Beaudry, P., D. Green and B. Sand, “The great reversal in the demand for skill and cognitive tasks”, *Journal of Labor Economics* **34**, S1, S199–S247 (2016).

Benhabib, J., R. Rogerson and R. Wright, “Homework in macroeconomics: Household production and aggregate fluctuations”, *Journal of Political Economy* **99**, 6, 1166–1187, URL <http://www.jstor.org/stable/2937726> (1991).

Bolt, J. and J. L. van Zanden, “The maddison project: Collaborative research on historical national accounts”, *Economic History Review* **67**, 3, 627–651, URL <http://dx.doi.org/10.1111/1468-0289.12032> (2014).

Bridgman, B., “Home productivity”, *Journal of Economic Dynamics and Control* pp. –, URL <http://www.sciencedirect.com/science/article/pii/S0165188916301348> (2016).

Buera, F. and J. Kaboski, “The rise of the service economy”, *The American Economic Review* **102**, 6, 2540–2569, URL <http://login.ezproxy1.lib.asu.edu/login?url=http://search.proquest.com/docview/1082> (2012).

Buera, F. J. and J. P. Kaboski, “Can traditional theories of structural change fit the data?”, *Journal of the European Economic Association* **7**, 2-3, 469–477, URL <http://dx.doi.org/10.1162/JEEA.2009.7.2-3.469> (2009).

Buera, F. J., J. P. Kaboski and K. Zhao, “Quantifying the growth in services: the role of skills, scale, and female labor supply”, *Journal of Human Capital* (forthcoming).

Bárány, Z. L. and C. Siegel, “Job polarization and structural change”, *American Economic Journal: Macroeconomics* **10**, 1, 57–89, URL <http://www.aeaweb.org/articles?id=10.1257/mac.20150258> (2018).

Carrillo-Tudela, C., B. Hobijn, P. She and L. Visschers, “The extent and cyclicity of career changes: Evidence for the u.k.”, *European Economic Review* **84**, 18 – 41, URL <http://www.sciencedirect.com/science/article/pii/S0014292115001464> (2016).

Carter, S. B., S. S. Gartner, M. R. Haines, A. L. Olmstead, R. Sutch and G. Wright, *The Historical Statistics of the United States Online Edition: Millennial Edition* (Cambridge University Press, New York, 2006), URL <https://books.google.com/books?id=4FqQNAEACAAJ>.

Caselli, F., “Chapter 9 accounting for cross-country income differences”, vol. 1 of *Handbook of Economic Growth*, pp. 679 – 741 (Elsevier, 2005), URL <http://www.sciencedirect.com/science/article/pii/S1574068405010099>.

Caselli, F., “Experience-biased technical change”, Discussion Papers 1517, Centre for Macroeconomics (CFM), URL <https://ideas.repec.org/p/cfm/wpaper/1517.html> (2015).

- Caselli, F. and J. Coleman, “The U.S. Structural Transformation and Regional Convergence: A Reinterpretation”, *Journal of Political Economy* **109**, 3, 584–616, URL <https://ideas.repec.org/a/ucp/jpolec/v109y2001i3p584-616.html> (2001a).
- Caselli, F. and W. J. Coleman, “The u.s. structural transformation and regional convergence: A reinterpretation”, *Journal of Political Economy* **109**, 3, 584–616 (2001b).
- Caselli, F., G. Esquivel and F. Lefort, “Reopening the convergence debate: A new look at cross-country growth empirics”, *Journal of Economic Growth* **1**, 3, 363–389, URL <http://www.jstor.org/stable/40215922> (1996).
- Cerina, F., A. Moro and M. P. Rendall, “The Role of Gender in Employment Polarization”, Discussion Papers 1704, Centre for Macroeconomics (CFM), URL <https://ideas.repec.org/p/cfm/wpaper/1704.html> (2017).
- Cociuba, S. and J. C. MacGee, “Demographics and sectoral reallocations: A search theory with immobile workers”, Mimeo, University of Western Ontario (2018).
- Cortés, G. M., “Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data”, *Journal of Labor Economics* **34**, 1, 63 – 105, URL <https://ideas.repec.org/a/ucp/jlabec/doi10.1086-682289.html> (2016).
- Cortés, G. M., N. Jaimovich, C. Nekarda and H. Siu, “The Micro and Macro of Disappearing Routine Jobs: A Flows Approach”, NBER Working Papers 20307, National Bureau of Economic Research, Inc, URL <https://ideas.repec.org/p/nbr/nberwo/20307.html> (2014).
- Duarte, M. and D. Restuccia, “The role of the structural transformation in aggregate productivity”, *The Quarterly Journal of Economics* **125**, 1, 129–173, URL <http://www.jstor.org/stable/40506279> (2010).
- Duarte, M. and D. Restuccia, “Relative prices and sectoral productivity”, NBER Working Paper 23979 (2017).
- Duernecker, G. and B. Herrendorf, “Structural transformation of occupation employment”, Working paper, University of Mannheim & Arizona State University (2016).
- Duernecker, G. and B. Herrendorf, “Structural transformation of occupation employment”, Mimeo, Arizona State University (2017).
- Echevarria, C., “Changes in sectoral composition associated with economic growth”, *International Economic Review* **38**, 2, 431–452 (1997).
- Fair, R. C. and J. B. Taylor, “Solution and Maximum Likelihood Estimation of Dynamic Nonlinear Rational Expectations Models”, *Econometrica* **51**, 4, 1169–1185, URL <https://ideas.repec.org/a/ecm/emetrp/v51y1983i4p1169-85.html> (1983).

- Fernández, R., A. Fogli and C. Olivetti, “Mothers and sons: Preference formation and female labor force dynamics”, *The Quarterly Journal of Economics* **119**, 4, 1249–1299, URL <http://qje.oxfordjournals.org/content/119/4/1249.abstract> (2004).
- Firpo, S., N. Fortin and T. Lemieux, “Occupational Tasks and Changes in the Wage Structure”, IZA Discussion Papers 5542, Institute for the Study of Labor (IZA), URL <https://ideas.repec.org/p/iza/izadps/dp5542.html> (2011).
- Freeman, R. and R. Schettkat, “Marketization of household production and the eu-us gap in work”, *Economic Policy* **20**, 41, 7–50, URL <http://www.jstor.org/stable/3601063> (2005).
- Gaffield, C., P. Baskerville, S. Cadigan, M. St-Hilaire, C. Bellavance, F. Normand, G. Darroch, C. Amrhein, L. Tepperman, C. Jones and E. Sager, “National sample of the 1911 census of canada [dataset]”, Edmonton, Canada: University of Alberta [distributor] (2009).
- Goldin, C., “The Quiet Revolution That Transformed Women’s Employment, Education, and Family”, *American Economic Review* **96**, 2, 1–21, URL <https://ideas.repec.org/a/aea/aecrev/v96y2006i2p1-21.html> (2006).
- Goos, M. and A. Manning, “Lousy and Lovely Jobs: The Rising Polarization of Work in Britain”, *The Review of Economics and Statistics* **89**, 1, 118–133, URL <https://ideas.repec.org/a/tpr/restat/v89y2007i1p118-133.html> (2007).
- Goos, M., A. Manning and A. Salomons, “Explaining job polarization: Routine-biased technological change and offshoring”, *American Economic Review* **104**, 8, 2509–26, URL <http://www.aeaweb.org/articles?id=10.1257/aer.104.8.2509> (2014).
- Greenwood, J., Z. Hercowitz and P. Krusell, “Long-Run Implications of Investment-Specific Technological Change”, *American Economic Review* **87**, 3, 342–362, URL <https://ideas.repec.org/a/aea/aecrev/v87y1997i3p342-62.html> (1997).
- Greenwood, J. and M. Yorukoglu, “1974”, *Carnegie-Rochester Conference Series on Public Policy* **46**, 1, 49–95, URL <https://ideas.repec.org/a/eee/crcspp/v46y1997ip49-95.html> (1997).
- Hendricks, L., “Cross-country variation in educational attainment: Structural change or within industry skill upgrading?”, *Journal of Economic Growth* **15**, 3, 205–233 (2010).
- Herrendorf, B., R. Rogerson and A. Valentinyi, “Two Perspectives on Preferences and Structural Transformation”, *American Economic Review* **103**, 7, 2752–89, URL <https://ideas.repec.org/a/aea/aecrev/v103y2013i7p2752-89.html> (2013).
- Herrendorf, B., R. Rogerson and A. Valentinyi, “Growth and Structural Transformation”, in “Handbook of Economic Growth”, vol. 2 of *Handbook of Economic Growth*, chap. 6, pp. 855–941 (Elsevier, 2014a), URL <https://ideas.repec.org/h/eee/grochp/2-855.html>.

- Herrendorf, B., R. Rogerson and A. Valentinyi, “Growth and structural transformation”, in “Handbook of Economic Growth”, edited by P. Aghion and S. N. Durlauf, vol. 2 (Elsevier, 2014b).
- Herrendorf, B. and T. Schoellman, “Wages, human capital, and barriers to structural transformation”, *American Economic Journal: Macroeconomics* **10**, 2, 1–23 (2018).
- Hill, M., “Patterns of time use”, in “Time, Goods, and Well-Being”, edited by T. Juster and F. Stafford (University of Michigan, 1985).
- Hobijn, B., F. Ravenna and A. Tambalotti, “Menu Costs at Work: Restaurant Prices and the Introduction of the Euro”, *Quarterly Journal of Economics* **121**, 3, 1103–1131, URL <https://ideas.repec.org/a/oup/qjecon/v121y2006i3p1103-1131..html> (2006).
- Hsieh, C.-T. and P. J. Klenow, “Development accounting”, *American Economic Journal: Macroeconomics* **2**, 1, 207–23, URL <http://www.aeaweb.org/articles?id=10.1257/mac.2.1.207> (2010).
- Inwood, K. and J. Chelsea, “National sample of the 1891 census of canada”, Guelph, Canada: University of Guelph (2011).
- Jaimovich, N. and H. Siu, “The trend is the cycle: Job polarization and jobless recoveries”, Working Paper 18334, National Bureau of Economic Research, URL <http://www.nber.org/papers/w18334> (2012).
- Juhn, C. and S. Potter, “Changes in labor force participation in the united states”, *Journal of Economic Perspectives* **20**, 3, 27–46, URL <http://EconPapers.repec.org/RePEc:aea:jecper:v:20:y:2006:i:3:p:27-46> (2006).
- Kambourov, G. and I. Manovskii, “Rising occupational and industry mobility in the united states: 1968-97”, *International Economic Review* **49**, 1, 41–79 (2008).
- Kambourov, G. and I. Manovskii, “Occupational Specificity Of Human Capital”, *International Economic Review* **50**, 1, 63–115, URL <https://ideas.repec.org/a/ier/iecrev/v50y2009i1p63-115.html> (2009).
- Katz, L. and K. Murphy, “Changes in relative wages, 1963–1987: Supply and demand factors”, *The Quarterly Journal of Economics* **107**, 1, 35–78, URL <https://ideas.repec.org/a/oup/qjecon/v107y1992i1p35-78..html> (1992).
- Kendrick, J. W., *Productivity Trends in the United States* (Princeton University Press, Princeton, 1961), URL <https://EconPapers.repec.org/RePEc:nbr:nberbk:kend61-1>.
- Kim, D.-I. and R. H. Topel, “Labor markets and economic growth: Lessons from korea’s industrialization, 1970-1990”, in “Differences and Changes in Wage Structures”, pp. 227–264 (University of Chicago Press, 1995).

- Kongsamut, P., S. Rebelo and D. Xie, “Beyond balanced growth”, *The Review of Economic Studies* **68**, 4, 869–882, URL <http://restud.oxfordjournals.org/content/68/4/869.abstract> (2001a).
- Kongsamut, P., S. Rebelo and D. Xie, “Beyond balanced growth”, *Review of Economic Studies* **68**, 4, 869–882 (2001b).
- Kuznets, S., *National Income and Its Composition, 1919-1938* (National Bureau of Economic Research, Inc, 1941), URL <http://EconPapers.repec.org/RePEc:nbr:nberbk:kuzn41-1>, with Lillian Epstein & Elizabeth Jenks.
- Kuznets, S., “Modern economic growth: Findings and reflections”, *The American Economic Review* **63**, 3, 247–258, URL <http://www.jstor.org/stable/1914358> (1973).
- Lagakos, D. and M. Waugh, “Specialization, agriculture, and cross-country productivity differences”, *American Economic Review* **103**, 2, 948–980 (2013).
- Mandelman, F., “Labor market polarization and international macroeconomic dynamics”, *Journal of Monetary Economics* **79**, 1 – 16, URL <http://www.sciencedirect.com/science/article/pii/S0304393216000076> (2016).
- Menzio, G., I. Telyukova and L. Visschers, “Directed Search over the Life Cycle”, *Review of Economic Dynamics* **19**, 38–62, URL <https://ideas.repec.org/a/red/issued/14-237.html> (2016).
- Minnesota Population Center, “Integrated public use microdata series — international: Version 6.3 [machine-readable database].”, Minneapolis: University of Minnesota., <http://www.ipums.org> (2014).
- Minnesota Population Center, “North atlantic population project: Complete count microdata.”, Minneapolis: Minnesota Population Center, available online at <http://doi.org/10.18128/D040.V2.3>. (2017).
- Minnesota Population Center, “Integrated public use microdata series, international: Version 7.1 [dataset]”, Minneapolis, MN: IPUMS (2018).
- Ngai, L. R. and B. Petrongolo, “Gender gaps and the rise of the service economy”, *American Economic Journal: Macroeconomics* **9**, 4, 1–44 (2017).
- Ngai, L. R. and C. A. Pissarides, “Structural change in a multisector model of growth”, *American Economic Review* **97**, 1, 429–443 (2007a).
- Ngai, L. R. and C. A. Pissarides, “Trends in Hours and Economic Growth”, *Review of Economic Dynamics* **11**, 2, 239–256, URL <https://ideas.repec.org/a/red/issued/07-15.html> (2008a).

- Ngai, R. and B. Petrongolo, “Gender gaps and the rise of the service economy”, Discussion Papers 1404, Centre for Macroeconomics (CFM), URL <http://EconPapers.repec.org/RePEc:cfm:wpaper:1404> (2014).
- Ngai, R. and C. Pissarides, “Structural Change in a Multisector Model of Growth”, *American Economic Review* **97**, 1, 429–443, URL <https://ideas.repec.org/a/aea/aecrev/v97y2007i1p429-443.html> (2007b).
- Ngai, R. and C. Pissarides, “Trends in Hours and Economic Growth”, *Review of Economic Dynamics* **11**, 2, 239–256, URL <https://ideas.repec.org/a/red/issued/07-15.html> (2008b).
- Parente, S. L. and E. C. Prescott, “Barriers to technology adoption and development”, *Journal of Political Economy* **102**, 2, 298–321 (1994).
- Perez, S., “The impact of railroads on education and occupational attainment: Evidence from 19th century argentina”, Mimeo, Stanford University (2016).
- Porzio, T. and G. Santangelo, “Structural change and the supply of agricultural workers”, Mimeo, University of California - San Diego (2017).
- Prescott, E., “Why do americans work so much more than europeans?”, *Quarterly Review* **28**, 1, 2–13, URL <https://ideas.repec.org/a/fip/fedmqr/y2004ijulp2-13nv.28no.1.html> (2004).
- Rendall, M., “Brain versus brawn: the realization of women’s comparative advantage”, Mimeo, Monash University (2017).
- Rogerson, R., “Market work, home work, and taxes: A cross-country analysis”, *Review of International Economics* **17**, 3, 588–601, URL <https://ideas.repec.org/a/bla/reviec/v17y2009i3p588-601.html> (2009).
- Ruggles, S., T. J. Alexander, K. Genadek, R. Goeken, M. B. Schoreder and M. Sobek, “Integrated public use microdata series: Version 5.0 [machine-readable database]”, Minneapolis, MN: Minnesota Population Center [producer and distributor], available online at <http://usa.ipums.org/usa/>. (2010).
- Ruggles, S., K. Genadek, R. Goeken, J. Grover and M. Sobek, “Integrated public use microdata series: Version 6.0”, [machine-readable database], Minneapolis: University of Minnesota (2015).
- Sala-i-Martin, X., “Cross-sectional regressions and the empirics of economic growth”, *European Economic Review* **38**, 3, 739 – 747, URL <http://www.sciencedirect.com/science/article/pii/0014292194901090> (1994).
- Schultz, T. W., *The Economic Organization of Agriculture* (McGraw-Hill New York, New York, 1953).

- The Digital Archive (The National Archive), Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center, “National sample of the 1865 census of norway, version 2.0”, Tromsø, Norway: University of Tromsø (2008a).
- The Digital Archive (The National Archive), Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center, “National sample of the 1900 census of norway, version 2.0”, Tromsø, Norway: University of Tromsø (2008b).
- The Digital Archive (The National Archive), Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center, “National sample of the 1910 census of norway, version 1.0”, Tromsø, Norway: University of Tromsø (2011).
- The Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center, “National sample of the 1875 census of norway, version 2.0”, Tromsø, Norway: University of Tromsø (2008).
- Tüzemen, D. and J. Willis, “The vanishing middle: Job polarization and workers’ response to the decline in middle-skill jobs”, *Economic Review* **98**, 1, 5–32, URL <https://ideas.repec.org/a/fip/fedker/y2013iqip5-32nv.98no.1.html> (2013).
- Vindas, A., “Job polarization, structural transformation and labor force participation”, Working paper, Arizona State University (2017).
- Young, A., “Structural transformation, the mismeasurement of productivity growth, and the cost disease of services”, *American Economic Review* **104**, 11, 3635–67, URL <http://www.aeaweb.org/articles?id=10.1257/aer.104.11.3635> (2014).

APPENDIX A
APPENDICES TO CHAPTER 1

A.1 Data Sources

The data I use covers the 1968-2018 period, and comes from the employment data in the Annual Social and Economic (ASEC) supplement to the Current Population Survey. I accessed these databases from the IPUMS-CPS project, an integrated set of data from the Current Population Surveys that goes through a convenient harmonization process. I consider the population aged between 25 and 65 years, and use their labels to determine industry, occupation (for the employed), and labor force status.

I construct the industry categories by grouping into goods and services. The goods industry includes both the manufacturing and agriculture sectors, which encompass manufacturing, construction, mining, agriculture, forestry, and fishery industrial classifications. The industrial classifications of the services industry are transportation, communications, public utilities, wholesale trade, retail trade, finance, insurance, real estate, business and repair services, personal services, entertainment and recreation services, professional and related services, and public administration.

To construct the occupation categories, I follow Cortés *et al.* (2014). They classify occupations based on two criteria: whether the tasks they involve are primarily manual or cognitive, and whether these are of a routine nature or not. Jaimovich and Siu (2012) note that there is a ranking in terms of wages: non-routine cognitive earn the highest while non-routine manual the lowest. In terms of tasks, non-routine cognitive tend to be high-skilled, while non-routine manual tend to be low-skilled. Routine manual and routine cognitive tend to be middle-skilled, so I group them together, in a similar fashion to Cortés (2016). I end up with three occupation groups then: non-routine manual, routine, and non-routine cognitive. Due to their association with the skills required, in the rest of the article I refer to these as manual, routine, and abstract occupations.

With every decennial Census, the occupation classifications are revised. These imply discrete jumps in their structure; even with the coarse grouping I use the changes are visible. The biggest changes were made with the 1983 and 2003 Censuses, which explain some of the shifts in the figures presented later. Both the harmonization processes from the IPUMS project and the analyses in Cortés *et al.* (2014) are careful enough to try and minimize these effects. In terms of the longer time trends, these reclassifications do not alter overall patterns, and do not represent a significant concern.

A.2 Detailed Tables

Table A.1: Occupational Job Polarization

	1968	2018
Manual	11.2	14.2
Routine	59.3	42.9
Abstract	29.5	42.9

These percentages refer to each occupation's share in employment.

Source: author's calculations using CPS.

Table A.2: Job Polarization and Non-employment

	1968	2018
Manual	7.6	11.1
Routine	40.0	33.7
Abstract	19.9	33.6
Non-employment	32.6	21.6

These percentages refer to each category's share in the total population.

Source: author's calculations using CPS.

Table A.3: Occupation Shares within Industries

	Goods		Services	
	1968	2018	1968	2018
Manual	1.2	1.3	18.1	17.6
Routine	76.7	68.5	47.4	36.2
Abstract	22.1	30.2	34.5	46.2

These percentages refer to the share of each occupation the industry's labor demand.

Source: author's calculations using CPS.

Table A.4: Industry Shares

	1968	2018
Goods	40.6	20.9
Services	59.4	79.1

These percentages refer to the share of each industry in the labor force.

Source: author's calculations using CPS.

A.3 Model

In this section, I present a static model of labor allocation between occupations to study the patterns shown earlier. It is a variation of Duernecker and Herrendorf (2016) and Ngai and Petrongolo (2014): a model of structural transformation, that features occupational choices within the firms, and allows for labor non-participation by including a home production sector.

The agents in this model choose between market and non-market work. In the market, firms decide how to allocate their labor into the three market occupations: manual, routine or abstract. In non-market work, agents devote time exclusively to home production. The driving force is *occupation specific* technical progress, and the difference in their growth rates induces the three main results: polarization, structural transformation, and changes in the labor force participation.

A.3.1 Environment

This is a discrete-time model where time runs forever. On the production side, I follow Ngai and Petrongolo (2014) and study three productive sectors: goods, market services, and home services. To distinguish between the jobs agents are working in, and the industries where these take place, I denote by lowercase j the occupation, and by uppercase I the industry. Then, $j \in \{h, m, r, a\}$, meaning these jobs can be in home production, manual occupations, routine occupations, and abstract occupations. Similarly, $I \in \{G, M, H\}$ denotes the production of goods, of market services, and of home services.

Home services are produced with a linear technology on home labor:

$$Y_{Ht} = A_{ht}N_{Hht} \quad (\text{A.1})$$

where A_{ht} denotes the efficiency of home production, and N_{Hht} denotes the amount of labor used in home production.

The firms in goods and market services produce with a technology that requires the three types of market occupations: manual, routine, and abstract. These are combined according to a CES aggregator:

$$Y_{It} = \left[\sum_{j \in \{m, r, a\}} \alpha_{Ij}^{\frac{1}{\sigma}} (A_{jt}N_{Ijt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \quad (\text{A.2})$$

where $I \in \{G, M\}$ indicates the industry and $j \in \{m, r, a\}$ the occupation. In this setting, N_{Ijt} denotes the input that industry I uses of occupation j , and A_{jt} is the labor efficiency in occupation j . The elasticity of substitution between the labor inputs is $\sigma > 0$, and $\alpha_{Ij} \in (0, 1)$ is the intensity of occupation j in sector I . This productive structure is similar to Duernecker and Herrendorf (2016), where labor efficiency is occupation-specific as opposed to industry specific, which is the standard assumption in the structural transformation literature.

For notational convenience, define the following:

$$N_{It} = N_{Imt} + N_{Irt} + N_{Iat} \quad I \in \{G, M\} \quad (\text{A.3})$$

$$N_{jt} = N_{Gjt} + N_{Mjt} \quad j \in \{m, r, a\} \quad (\text{A.4})$$

$$\begin{aligned} N_t &= N_{mt} + N_{rt} + N_{at} \\ &= N_{Gt} + N_{Mt} \end{aligned} \quad (\text{A.5})$$

where N_{It} is the total amount of labor in industry $I \in \{G, M\}$, N_{jt} is the total amount of labor in occupation $j \in \{m, r, a\}$, and N_t is total market labor, which is clearly equal to the sum of labor over market industries or occupations.

On the consumption side, there are identical households of measure one. These consume goods and a combination of home and market services. The utility level they yield is aggregated according to a nested CES specification:

$$U_t(C_{Gt}, C_{St}) = \left[\omega_G^{\frac{1}{\varepsilon}} (C_{Gt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_S^{\frac{1}{\varepsilon}} (C_{St})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \quad (\text{A.6})$$

where C_{Gt} and C_{St} denote the consumption of goods and compound services. Their relative weights are ω_G and ω_S , which add up to one, and individually are between zero and one. The elasticity of substitution between goods and compound services is $\varepsilon > 0$.

Compound services are also aggregated through a CES specification:

$$C_{St} = \left[\varphi_M^{\frac{1}{\eta}} (C_{Mt})^{\frac{\eta-1}{\eta}} + \varphi_H^{\frac{1}{\eta}} (C_{Ht})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} \quad (\text{A.7})$$

where C_{Mt} and C_{Ht} denote market and home services. Their relative weights are φ_M and φ_H , which also add up to one, and individually are between zero and one. The elasticity of substitution between market and home services is $\eta > 0$.

All households are endowed with one unit of labor in each period. I denote by L_t total labor supply, and the remaining $1 - L_t$ is devoted to home production. Within the market, labor is perfectly mobile across occupations, and has no occupation or sector specificity to it.

The feasibility conditions for the consumption sectors are:

$$Y_{Gt} = C_{Gt} \quad (\text{A.8})$$

$$Y_{Mt} = C_{Mt} \quad (\text{A.9})$$

$$Y_{Ht} = C_{Ht} \quad (\text{A.10})$$

These equations simply require that what is produced in the goods, market services and home services industries be consumed by the households.

Finally, the feasibility condition for the labor market requires that the households' labor supply be equal to the market demand:

$$\begin{aligned} L_t &= N_{mt} + N_{rt} + N_{at} \\ &= N_{Gt} + N_{Mt} \\ &= N_t \end{aligned} \quad (\text{A.11})$$

A.3.2 Decentralized Market Structure

I assume competitive markets for labor, goods and market services, where all agents take the prices as given. Firms in the consumption and market services industries face the following profit-maximization problem:

$$\max_{N_{Imt}, N_{Irt}, N_{Iat}} p_{It} \left[\sum_{j \in \{m, r, a\}} \alpha_{Ij}^{\frac{1}{\sigma}} (A_{jt} N_{Ijt})^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} - w_t (N_{Imt} + N_{Irt} + N_{Iat}) \quad (\text{A.12})$$

where p_{It} is the market price of their output, and w_{jt} the market wages. First order conditions imply:

$$\frac{N_{Imt}}{N_{Irt}} = \frac{\alpha_{Im}}{\alpha_{Ir}} \left(\frac{A_{rt}}{A_{mt}} \right)^{1-\sigma} \quad (\text{A.13})$$

$$\frac{N_{Iat}}{N_{Imt}} = \frac{\alpha_{Ia}}{\alpha_{Im}} \left(\frac{A_{mt}}{A_{at}} \right)^{1-\sigma} \quad (\text{A.14})$$

Equations (A.13) and (A.14) describe the relative labor allocations between manual and routine occupations, and abstract and manual occupations, each for industry $I \in \{G, M\}$.

Taking prices as given, the household's utility-maximization problem at time t is:

$$\max_{L_t, C_{Gt}, C_{Mt}} \left[\omega_G^{\frac{1}{\varepsilon}} (C_{Gt})^{\frac{\varepsilon-1}{\varepsilon}} + \omega_S^{\frac{1}{\varepsilon}} (C_{St})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \quad (\text{A.15})$$

subject to:

$$\begin{aligned} p_{Gt} C_{Gt} + p_{Mt} C_{Mt} &= w_t L_t \\ C_{St} &= \left[\varphi_M^{\frac{1}{\eta}} (C_{Mt})^{\frac{\eta-1}{\eta}} + \varphi_H^{\frac{1}{\eta}} (C_{Ht})^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}} \\ C_{Ht} &= A_{ht} (1 - L_t) \end{aligned}$$

Households then, maximize their utility subject to their budget constraint in market products, and their technology constraint in household production. This problem can be solved in two steps: the first is to find the optimal allocation between home and market services, and the second one is for the optimal allocation between goods and compound services.

Home services are not traded in the market, meaning there is no market price attached to them. Its opportunity cost, however, is well defined since its alternatives have market prices attached to them. I denote by p_{Ht} this implicit price. The first order conditions to maximize C_{St} imply:

$$\frac{p_{Mt} C_{Mt}}{p_{Ht} C_{Ht}} = \frac{\varphi_M}{\varphi_H} \left(\frac{p_{Mt}}{p_{Ht}} \right)^{1-\eta} \quad (\text{A.16})$$

Define the following price index:

$$p_{St} = [\varphi_H(p_{Ht})^{1-\eta} + \varphi_M(p_{Mt})^{1-\eta}]^{\frac{1}{1-\eta}} \quad (\text{A.17})$$

This price index can be interpreted as the unit price of the optimal services basket, which is relevant for the decision between the consumption of goods and the composite services basket. For this decision, first order conditions imply:

$$\frac{p_{St}C_{St}}{p_{Gt}C_{Gt}} = \frac{\omega_S}{\omega_G} \left(\frac{p_{St}}{p_{Gt}} \right)^{1-\varepsilon} \quad (\text{A.18})$$

A.4 Estimation Procedure

In this section, I explain in further detail how to match the data to the model's parameters. Recall the assumption of constant growth rates in productivity; because of that I only need to look at the initial and final years. These are denoted by $t = 0$ and $t = T$.

To get the market occupation intensities in production (α_{Ij}), I use equations (A.13) (A.14). For the initial year, these imply:

$$\alpha_{Ir} = \alpha_{Im} \frac{N_{Ir0}}{N_{Im0}} \quad \alpha_{Ia} = \alpha_{Im} \frac{N_{Ia0}}{N_{Im0}} \quad (\text{A.19})$$

These two yield the intensities, since all three add up to one.

To get the relative weights in the consumption of market and home services (φ_H and φ_M), and in the consumption of goods and services (ω_G and ω_S), I use equations (1.13) and (1.15) in a similar fashion:

$$\varphi_H = \varphi_M \frac{N_{H0}}{N_{M0}} \quad \omega_G = \omega_S \varphi_M \frac{N_{G0}}{N_{M0}} \quad (\text{A.20})$$

To get the elasticity of substitution in the production function, I first rewrite the average labor productivity (1.10):

$$\begin{aligned} \tilde{A}_{It} &= A_{rt} \left\{ \alpha_{Ir} \left[1 + \frac{\alpha_{Im}}{\alpha_{Ir}} \left(\frac{A_{rt}}{A_{mt}} \right)^{1-\sigma} + \frac{\alpha_{Ia}}{\alpha_{Ir}} \left(\frac{A_{rt}}{A_{at}} \right)^{1-\sigma} \right] \right\}^{\frac{-1}{1-\sigma}} \\ &= A_{rt} \left\{ \alpha_{Ir} \left[1 + \frac{N_{Imt}}{N_{Irt}} + \frac{N_{Iat}}{N_{Irt}} \right] \right\}^{\frac{-1}{1-\sigma}} \end{aligned} \quad (\text{A.21})$$

Then, I use equation (1.15) in the final period, substituting in equations (1.13) and (A.21):

$$\frac{N_{MT}}{N_{GT}} = \frac{\omega_S}{\omega_G} \left[\frac{\alpha_{Mr} \left(1 + \frac{N_{MmT}}{N_{MrT}} + \frac{N_{MaT}}{N_{MrT}} \right)}{\alpha_{Gr} \left(1 + \frac{N_{GmT}}{N_{GrT}} + \frac{N_{GaT}}{N_{GrT}} \right)} \right]^{\frac{1-\varepsilon}{1-\sigma}} \left(\varphi_M \frac{N_{ST}}{N_{MT}} \right)^{\frac{1-\varepsilon}{1-\eta}} \quad (\text{A.22})$$

Applying logarithms and rearranging leads to my estimate of σ .

To get relative occupational productivities, I use the occupation shares in period T , and equations (A.13) and (A.14) for the market services industry. Rearranging, these give:

$$\frac{A_{mT}}{A_{rT}} = \left(\frac{\alpha_{Mm} N_{MrT}}{\alpha_{Mr} N_{MmT}} \right)^{\frac{1}{1-\sigma}} \quad \frac{A_{aT}}{A_{rT}} = \left(\frac{\alpha_{Ma} N_{MrT}}{\alpha_{Mr} N_{MaT}} \right)^{\frac{1}{1-\sigma}} \quad (\text{A.23})$$

To get the productivity levels in market occupations, I use data from the Bureau of Economic Analysis. In particular, I take the real gross domestic product per capita time series (A939RX0Q048SBEA) to establish that this has grown by a factor of 2.23 between 1968 and 2018. To reproduce this growth pattern, I match this factor with market production evaluated at year 0's prices. Substituting (A.21) into the production function (A.2):

$$\begin{aligned} 2.23 &= \frac{p_{M0}Y_{MT} + p_{G0}Y_{GT}}{p_{M0}Y_{M0} + p_{G0}Y_{G0}} \\ &= \frac{A_{rT} N_{MT} \left[\alpha_{Mr} \left(1 + \frac{N_{MmT}}{N_{MrT}} + \frac{N_{MaT}}{N_{MrT}} \right) \right]^{\frac{-1}{1-\sigma}} + N_{GT} \left[\alpha_{Gr} \left(1 + \frac{N_{GmT}}{N_{GrT}} + \frac{N_{GaT}}{N_{GrT}} \right) \right]^{\frac{-1}{1-\sigma}}}{A_{r0} N_{M0} \left[\alpha_{Mr} \left(1 + \frac{N_{Mm0}}{N_{Mr0}} + \frac{N_{Ma0}}{N_{Mr0}} \right) \right]^{\frac{-1}{1-\sigma}} + N_{G0} \left[\alpha_{Gr} \left(1 + \frac{N_{Gm0}}{N_{Gr0}} + \frac{N_{Ga0}}{N_{Gr0}} \right) \right]^{\frac{-1}{1-\sigma}}} \end{aligned} \quad (\text{A.24})$$

This yields the growth factor of productivity in routine occupations. With this I reconstruct the other productivity levels for market occupations.

Finally, from equation (1.13) in the final year, I get the productivity level in home production:

$$\frac{A_{rT}}{A_{hT}} = \left[\frac{\varphi_M N_{HT}}{\varphi_H N_{MT}} \right]^{\frac{1}{1-\eta}} \left[\alpha_{Mr} \left(1 + \frac{N_{MmT}}{N_{MrT}} + \frac{N_{MaT}}{N_{MrT}} \right) \right]^{\frac{1}{1-\sigma}} \quad (\text{A.25})$$

APPENDIX B
APPENDICES TO CHAPTER 2

B.1 United States Data

For the United States we draw on census microdata, taken from Ruggles *et al.* (2010). We exclude the 1850 and 1860 censuses because they did not enumerate slaves. The 1890 census microdata were lost in a fire. We pool the 2008–2012 American Community Surveys to take the place of a 2010 census, in line with usual IPUMS practice. The pooled surveys have very similar questions, responses, and coverage as the 2000 census. Together, we have 14 surveys spanning 140 years at ten year intervals, excluding 1890.

We focus throughout on workers who are employed and have a valid response to age and industry of employment. We aggregate industry codes to three broad industry groups. Agriculture includes farming, forestry, and fishing. Manufacturing includes most remaining goods production, including manufacturing, mining, and construction. Services includes the remaining industries: utilities; retail and wholesale trade; hotels and restaurants; finance, insurance, and real estate; public administration and defense; education; health and social services; private household services; and other/miscellaneous services. We discard reported industry of employment for those younger than 20, under the view that this likely represents part-time or seasonal work and not a serious career choice.

Our demographic analysis focuses on the role of education, gender, and marital status in accounting for structural transformation. We focus here on the years 1940–2010, since education is available in the United States only from 1940 onward. We focus on workers with valid responses to all three questions. We aggregate the detailed variables so that we have broader categories that are easily comparable over time. We have two gender categories (male and female); two marriage categories (married and unmarried, which includes separated, divorced, and widowed); and four education categories (less than primary complete, primary complete, secondary complete, and tertiary complete).

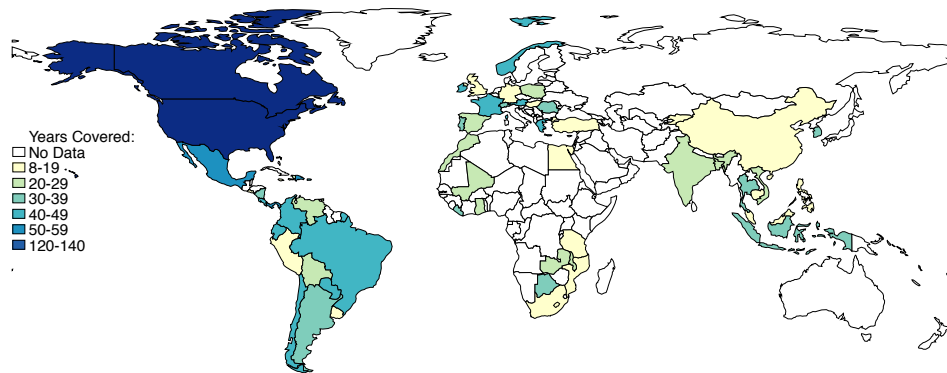
B.2 International Data

Our international sample includes all countries for which we have been able to acquire repeated (at least two) nationally representative cross-sections of microdata that include employment status and industry of employment. We further limit our attention data that include sufficient detail that we can harmonize these key objects in a reasonably consistent way; we also eliminate a few countries or samples that cover shorter periods (generally, less than eight years) to avoid confusing temporary changes with trends.

Most of our data come from IPUMS. Minnesota Population Center (2014) includes repeated cross-sections for 54 countries with the necessary information. Ruggles *et al.* (2010) provides the censuses for Puerto Rico. Minnesota Population Center (2017) includes additional cross-sections for Canada as well as new data for Norway from the late 19th and early 20th centuries.¹ Finally, we have identified three countries

¹Specific data provided by Inwood and Chelsea (2011), Gaffield *et al.* (2009), The Digital Archive (The National Archive), Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center (2008a), The Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center (2008), The Digital Archive (The National Archive), Norwegian His-

Figure B.1: Countries in Dataset



with independent data of sufficient information that we were able to harmonize and add to this dataset. All told, our dataset for studying worker reallocation includes 201 samples from 59 countries. Figure B.1 shows the countries covered and the total length spanned by country. Our coverage is broad both in terms of geography and PPP GDP per capita. Table B.1 includes a full list of countries and years.

The IPUMS team has devoted a great deal of energy to harmonizing variables and responses across countries and years. The most important for our purposes is that they have re-coded each country’s original responses for the industry or sector question (e.g., the one describing the activity or product produced at the respondent’s workplace) into a variable they call *indgen*, which is a slightly modified version of the ISIC 1-digit industry coding scheme. As they note, this coding process is non-trivial, in three main ways. First, for some countries the underlying codes are too coarse to be mapped into *indgen* at all; these countries are absent from our data. Second, in some countries not all of the original industry codes can be mapped into the *indgen* classification scheme. Finally, there are inevitably some judgment calls when constructing such crosswalks. The main examples described by the IPUMS team involve categories which are small (repair work) or judgment calls that are not relevant for our work (distinguishing among the service industries when mapping an industry). We aggregate these categories into the three broad industry groups as explained in the last subsection.

B.3 Alternative Industry Decompositions

In the text we focus on the classic three-industry description of structural transformation. This industry decomposition is useful for middle income countries, but less so for poor or rich countries, where agriculture or services respectively dominate. Here we show that our results are robust to studying alternative decompositions.

When studying poor countries, it is common to group all of non-agriculture together and focus simply on the transition from agriculture to non-agriculture. We

torical Data Centre (University of Tromsø) and the Minnesota Population Center (2008b), and The Digital Archive (The National Archive), Norwegian Historical Data Centre (University of Tromsø) and the Minnesota Population Center (2011).

Table B.1: Structural Transformation Sample

Country	N	Years	Country	N	Years
Argentina	4	1970–2001	Austria	5	1971–2011
Bangladesh	3	1991–2011	Bolivia	3	1976–2001
Botswana	4	1981–2011	Brazil	5	1970–2010
Cambodia	2	1998–2008	Canada	7	1891–2011
Chile	5	1960–2002	China	3	1982–2000
Colombia	4	1964–2005	Costa Rica	5	1963–2011
Dom. Republic	5	1960–2010	Ecuador	5	1962–2010
Egypt	2	1996–2006	El Salvador	2	1992–2007
Fiji	5	1966–2007	France	5	1962–2011
Germany (West)	2	1971–1981	Ghana	3	1984–2010
Greece	5	1971–2011	Haiti	2	1982–2003
Hungary	2	2001–2011	India	3	1983–2009
Indonesia	4	1971–2010	Ireland	4	1971–2011
Jamaica	3	1982–2001	Kyrgyzstan	2	1999–2009
Liberia	2	1974–2008	Malawi	3	1987–2008
Malaysia	2	1991–2000	Mali	3	1987–2009
Mexico	5	1960–2015	Morocco	3	1982–2004
Mozambique	2	1997–2007	Nicaragua	3	1971–2005
Norway	4	1865–1910	Palestine*	2	2000–2015
Panama	6	1960–2010	Paraguay	5	1962–2002
Peru	2	1993–2007	Philippines	2	1990–2000
Poland	2	1978–2002	Portugal	4	1981–2011
Puerto Rico	5	1970–2015	Romania	4	1977–2011
South Africa	2	1996–2007	South Korea*	4	1986–2016
Spain	3	1981–2001	Switzerland	4	1970–2000
Tanzania	2	2002–2012	Thailand	4	1970–2000
Trinidad & Tobago	3	1980–2000	Turkey	2	1985–2000
United Kingdom*	3	1997–2014	United States	14	1870–2015
Uruguay	2	1996–2006	Venezuela	3	1981–2001
Vietnam	3	1989–2009	Zambia	3	1990–2010

* Samples derived from independently collected labor force surveys. The remaining samples are from Ruggles *et al.* (2010), Minnesota Population Center (2014), and Minnesota Population Center (2017).

Table B.2: Between Cohort Share for Alternative Industry Groupings

Number of Broad Industries	Between Share
Two	71%
Four	49%
Fifteen	53%

showed in Table 2.1 that the between share is highest for agriculture. Thus, not surprisingly, this leads to a higher role for the between share in structural transformation, reported in the first row of Table B.2.

For rich countries it might be useful to decompose services, given that it now accounts for a large majority of employment (Duarte and Restuccia, 2017). We follow Herrendorf and Schoellman (2018) by breaking services into unskilled and skilled categories based on average education of the workforce. The former includes primarily personal services, wholesale and retail trade, and hotels and restaurants; the latter includes professional services. Figure B.2 shows graphically how structural transformation looks. Unskilled services display a mixed, possibly inverse-U shape similar to manufacturing, while skilled services grow uniformly. The second row of Table B.2 shows that this extra detail matters little for our basic metric; the between share in this case is 49 percent. To push this point even further, we can use a fifteen sector decomposition (the full detail of indgen as coded in ipums). With this decomposition the between share is still 53 percent. This suggests that the finding of an important role for new cohorts in accounting for structural transformation is robust.

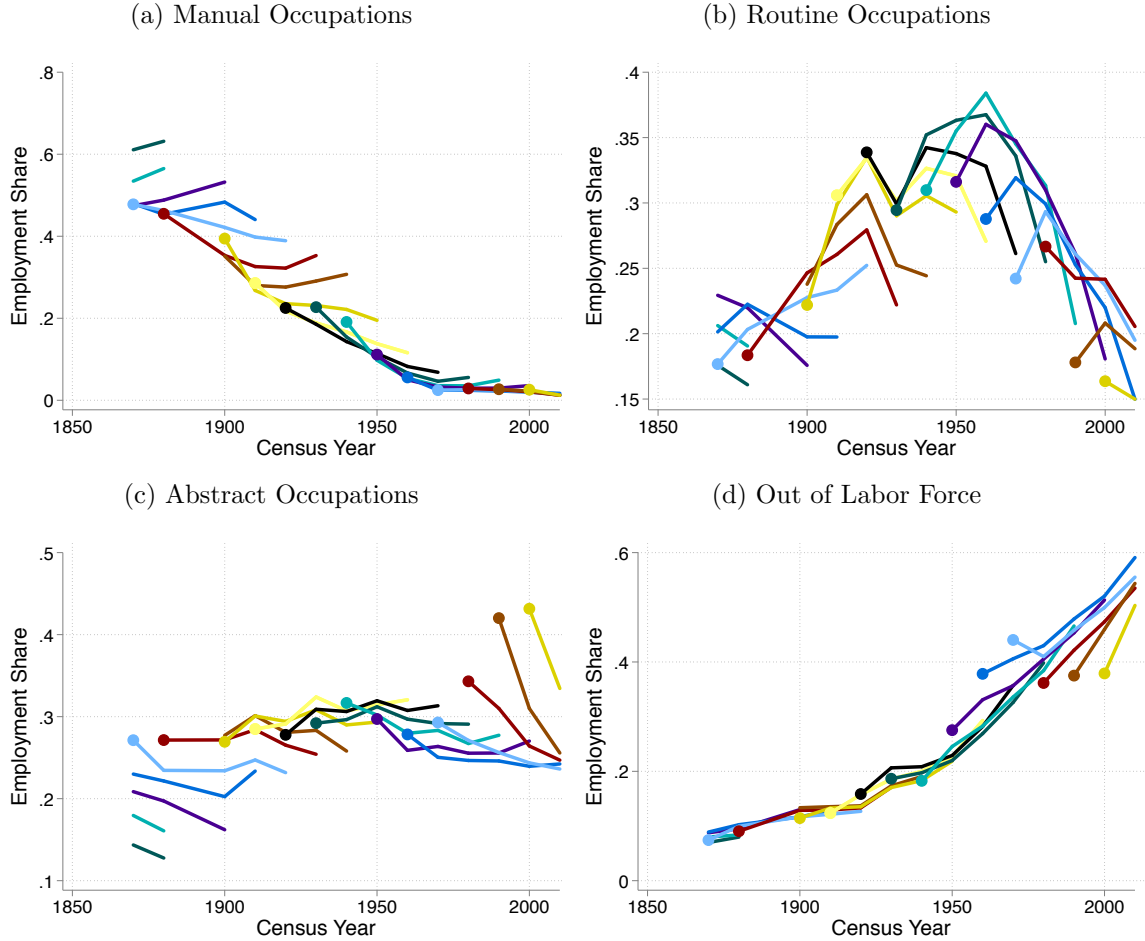
B.4 Solution to Career Choices

The solution to the dynamic discrete choices that make up the career decisions of workers in this model involves several conditional probabilities and conditional expectations of transformations of independently exponentially distributed random variables. In this section we derive these conditional probabilities and expectations in closed form. Because of the length of these derivations and the resulting expressions, they are omitted from the main text.

The variables we derive in closed form in this section are

- $\Phi_t(i)$, the probability of a worker of age $h = 0$ being trained to work in sector i at time t .
- $\tilde{z}_t(i; 0)$, the average productivity level of a worker trained to work in sector i at time t .
- $\Gamma_t(i, j; h)$, the probability that a worker of age $h > 0$ who works in sector i at time $t - 1$ decides to get retrained and starts working in sector j at time t .
- $\tilde{z}_t(i, j; h)$, the average productivity level of a worker of age $h > 0$ who decided to work in sector j at time t while having worked in sector i in period $t - 1$.

Figure B.2: Structural Transformation in the United States: 4 Industry View



All four of these variables can be derived using two results on the distribution and expectation of the maximum of three linear transformations of exponentially distributed random variables. We derive these results in general first below and then show how they apply to the four variables at hand.

Two main results about exponentially distributed random variables

The two results we consider are about linear transformations of three independently Exponentially distributed random variables, $Z_i \sim \text{Exp}(1)$ where $i = 1, \dots, 3$. We write these transformations as $X_i = a_i Z_i + b_i$, where $a_i, b_i > 0$.

The first result we use is that for the probability that X_i is the maximum of the sample of three X 's, which we denote by π_i . In terms of order statistics, this is

$$\pi_i = P[X_i = X_{(3)}] = P\left[a_i Z_i + b_i = \max_{j=1, \dots, 3} \{a_j Z_j + b_j\}\right]. \quad (\text{B.1})$$

Using the notation that $j \neq i$ and $k \neq j$ and $k \neq i$, we can write this probability as

$$\pi_i = \int_{z_i}^{\infty} \left(1 - \exp\left(-\frac{a_i Z + b_i - b_j}{a_j}\right)\right) \left(1 - \exp\left(-\frac{a_i Z + b_i - b_k}{a_k}\right)\right) \exp(-Z) dZ, \quad (B.2)$$

where

$$z_i = \max_{j=1,\dots,3} \left\{ \frac{b_j - b_i}{a_i} \right\}. \quad (B.3)$$

The above integral can be written in terms of four subintegrals. This yields

$$\pi_i = \int_{z_i}^{\infty} \exp(-Z) dZ \quad (B.4)$$

$$- \exp\left(\frac{b_j - b_i}{a_j}\right) \int_{z_i}^{\infty} \exp\left(-\left\{1 + \frac{a_i}{a_j}\right\} Z\right) dZ \quad (B.5)$$

$$- \exp\left(\frac{b_k - b_i}{a_k}\right) \int_{z_i}^{\infty} \exp\left(-\left\{1 + \frac{a_i}{a_k}\right\} Z\right) dZ \quad (B.6)$$

$$+ \exp\left(\frac{b_j - b_i}{a_j} + \frac{b_k - b_i}{a_k}\right) \int_{z_i}^{\infty} \exp\left(-\left\{1 + \frac{a_i}{a_k} + \frac{a_i}{a_j}\right\} Z\right) dZ \quad (B.7)$$

$$= \exp(-z_i) \quad (B.8)$$

$$- \exp\left(\frac{b_j - b_i}{a_j}\right) \left\{ \frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_j}} \right\} \exp\left(-\left\{1 + \frac{a_i}{a_j}\right\} z_i\right) \quad (B.9)$$

$$- \exp\left(\frac{b_k - b_i}{a_k}\right) \left\{ \frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_k}} \right\} \exp\left(-\left\{1 + \frac{a_i}{a_k}\right\} z_i\right) \quad (B.10)$$

$$+ \exp\left(\frac{b_j - b_i}{a_j} + \frac{b_k - b_i}{a_k}\right) \left\{ \frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_j} + \frac{1}{a_k}} \right\} \exp\left(-\left\{1 + \frac{a_i}{a_k} + \frac{a_i}{a_j}\right\} z_i\right) \quad (B.11)$$

For the numerical implementation, it is important to also take into account the limiting behavior of this integral when the a 's go to zero. If $a_j \downarrow 0$ then (B.5) and (B.7) are zero, and the integral only consists of (B.4) and (B.6). Similarly, if $a_k \downarrow 0$ then (B.6) and (B.7) are zero, and the integral only consists of (B.4) and (B.5).

Finally, if $a_i \downarrow 0$ then π_i is not an integral, but instead equals

$$\pi_i = \begin{cases} \left(1 - \exp\left(-\frac{b_i - b_j}{a_j}\right)\right) \left(1 - \exp\left(-\frac{b_i - b_k}{a_k}\right)\right) & \text{if } b_i > b_j \text{ and } b_i > b_k \\ 0 & \text{otherwise} \end{cases}. \quad (B.12)$$

In this expression, the first term goes to one when $a_j \downarrow 0$ and the second term goes to one when $a_k \downarrow 0$.

We summarize this definition by defining the function

$$\pi_i = \Pi(a_i, b_i, a_j, b_j, a_k, b_k), \quad (B.13)$$

which is what we will use later on to define two of the four variables of interest.

The second main result that we consider is about the expected value of Z_i conditional on X_i being the third order statistic. We denote this expectation by $\tilde{\zeta}_i$ and, formally, it is defined as

$$\zeta_i = \mathbb{E} [Z_i | X_i = X_{(3)}] = \mathbb{E} \left[Z_i \left| a_i Z_i + b_i = \max_{j=1, \dots, 3} \{a_j Z_j + b_j\} \right. \right]. \quad (\text{B.14})$$

Given this definition, we can write

$$\zeta_i = \frac{1}{\pi_i} \int_{z_i}^{\infty} Z \left(1 - \exp \left(-\frac{a_i Z + b_i - b_j}{a_j} \right) \right) \left(1 - \exp \left(-\frac{a_i Z + b_i - b_k}{a_k} \right) \right) \exp(-Z) dZ. \quad (\text{B.15})$$

This, again, can be written in terms of four subintegrals as

$$\zeta_i = \frac{1}{\pi_i} \int_{z_i}^{\infty} Z \exp(-Z) dZ \quad (\text{B.16})$$

$$- \frac{1}{\pi_i} \exp \left(\frac{b_j - b_i}{a_j} \right) \int_{z_i}^{\infty} Z \exp \left(- \left\{ 1 + \frac{a_i}{a_j} \right\} Z \right) dZ \quad (\text{B.17})$$

$$- \frac{1}{\pi_i} \exp \left(\frac{b_k - b_i}{a_k} \right) \int_{z_i}^{\infty} Z \exp \left(- \left\{ 1 + \frac{a_i}{a_k} \right\} Z \right) dZ \quad (\text{B.18})$$

$$+ \frac{1}{\pi_i} \exp \left(\frac{b_j - b_i}{a_j} + \frac{b_k - b_i}{a_k} \right) \int_{z_i}^{\infty} Z \exp \left(- \left\{ 1 + \frac{a_i}{a_k} + \frac{a_i}{a_j} \right\} Z \right) dZ. \quad (\text{B.19})$$

Each of these subintegrals can be solved by using the result that

$$\int x \exp(-bx) dx = -\frac{1}{b^2} (bx + 1) \exp(-bx). \quad (\text{B.20})$$

Doing so, we obtain that

$$\zeta_i = \frac{1}{\pi_i} (1 + z_i) \exp(-z_i) \quad (\text{B.21})$$

$$- \frac{1}{\pi_i} \exp \left(\frac{b_j - b_i}{a_j} \right) \left\{ \frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_j}} \right\}^2 \left(\left\{ 1 + \frac{a_i}{a_j} \right\} z_i + 1 \right) \exp \left(- \left\{ 1 + \frac{a_i}{a_j} \right\} z_i \right) \quad (\text{B.22})$$

$$- \frac{1}{\pi_i} \exp \left(\frac{b_k - b_i}{a_k} \right) \left\{ \frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_k}} \right\}^2 \left(\left\{ 1 + \frac{a_i}{a_k} \right\} z_i + 1 \right) \exp \left(- \left\{ 1 + \frac{a_i}{a_k} \right\} z_i \right) \quad (\text{B.23})$$

$$+ \frac{1}{\pi_i} \exp \left(\frac{b_j - b_i}{a_j} + \frac{b_k - b_i}{a_k} \right) \left\{ \frac{\frac{1}{a_i}}{\frac{1}{a_i} + \frac{1}{a_j} + \frac{1}{a_k}} \right\}^2 \times \quad (\text{B.24})$$

$$\left(\left\{ 1 + \frac{a_i}{a_j} + \frac{a_i}{a_k} \right\} z_i + 1 \right) \exp \left(- \left\{ 1 + \frac{a_i}{a_k} + \frac{a_i}{a_j} \right\} z_i \right).$$

Again, it is useful to consider the limiting cases. When $a_j \downarrow 0$ then (B.17) and (B.19) go to zero and the integral just consists of parts (B.16) and (B.18). When $a_k \downarrow 0$ then (B.18) and (B.19) go to zero and the integral consists of (B.16) and

(B.17). When $a_i \downarrow 0$ then the value of Z_i does not matter for the career choice. As a result, in that case $\zeta_i = \mathbb{E}[Z_i] = 1$.

We summarize this definition by defining the function

$$\zeta_i = \mathcal{Z}(a_i, b_i, a_j, b_j, a_k, b_k), \quad (\text{B.25})$$

which is what we will use later on to define two of the four variables of interest.

Career-choice probabilities and expected productivity levels

To shorten the notation, it turns out to be useful to define

$$\tilde{V}_{t+1}(i, h) = \frac{1 - \delta}{1 + r} \mathbb{E}_t V_{t+1}(i, h + 1; \mathbf{z}_{t+1}). \quad (\text{B.26})$$

Using the above definition and the results derived in the subsection above, we can now write

$$\Phi_t(i) = \Pi\left((1 - \phi) w_{i,t}, \tilde{V}_{t+1}(i, h), (1 - \phi) w_{j,t}, \tilde{V}_{t+1}(j, h), (1 - \phi) w_{k,t}, \tilde{V}_{t+1}(k, h)\right), \quad (\text{B.27})$$

as well as

$$\tilde{z}_t(i; 0) = \mathcal{Z}\left((1 - \phi) w_{i,t}, \tilde{V}_{t+1}(i, h), (1 - \phi) w_{j,t}, \tilde{V}_{t+1}(j, h), (1 - \phi) w_{k,t}, \tilde{V}_{t+1}(k, h)\right). \quad (\text{B.28})$$

Moreover, we can write

$$\Gamma_t(i, i; h) = \Pi\left(w_{i,t}, \tilde{V}_{t+1}(i, h), (1 - \gamma_h) w_{j,t}, \tilde{V}_{t+1}(j, h), (1 - \gamma_h) w_{k,t}, \tilde{V}_{t+1}(k, h)\right), \quad (\text{B.29})$$

and

$$\Gamma_t(i, j; h) = \Pi\left((1 - \gamma_h) w_{j,t}, \tilde{V}_{t+1}(j, h), w_{i,t}, \tilde{V}_{t+1}(i, h), (1 - \gamma_h) w_{k,t}, \tilde{V}_{t+1}(k, h)\right), \quad (\text{B.30})$$

when $j \neq i$.

The associated expected productivity levels are

$$\tilde{z}_t(i, i; h) = \mathcal{Z}\left(w_{i,t}, \tilde{V}_{t+1}(i, h), (1 - \gamma_h) w_{j,t}, \tilde{V}_{t+1}(j, h), (1 - \gamma_h) w_{k,t}, \tilde{V}_{t+1}(k, h)\right), \quad (\text{B.31})$$

and

$$\tilde{z}_t(i, j; h) = \mathcal{Z}\left((1 - \gamma_h) w_{j,t}, \tilde{V}_{t+1}(j, h), w_{i,t}, \tilde{V}_{t+1}(i, h), (1 - \gamma_h) w_{k,t}, \tilde{V}_{t+1}(k, h)\right), \quad (\text{B.32})$$

when $j \neq i$.

B.5 Calibration of Flexible Benchmark Parameters

For the dynamics of employment shares in the model, what matters is the product of the preference parameter, λ_i , and the technology parameter, A_i . For this reason, we normalize the preference parameters² to be equal across industries

$$\lambda_a = \lambda_m = \lambda_s = \frac{1}{3}. \quad (\text{B.33})$$

Moreover, because preferences are homothetic and the production technology exhibits constant returns to scale, the absolute levels of the technology parameters (A_i) do not matter for the equilibrium dynamics, but their relative levels. With that in mind, we also normalize

$$A_a = 1. \quad (\text{B.34})$$

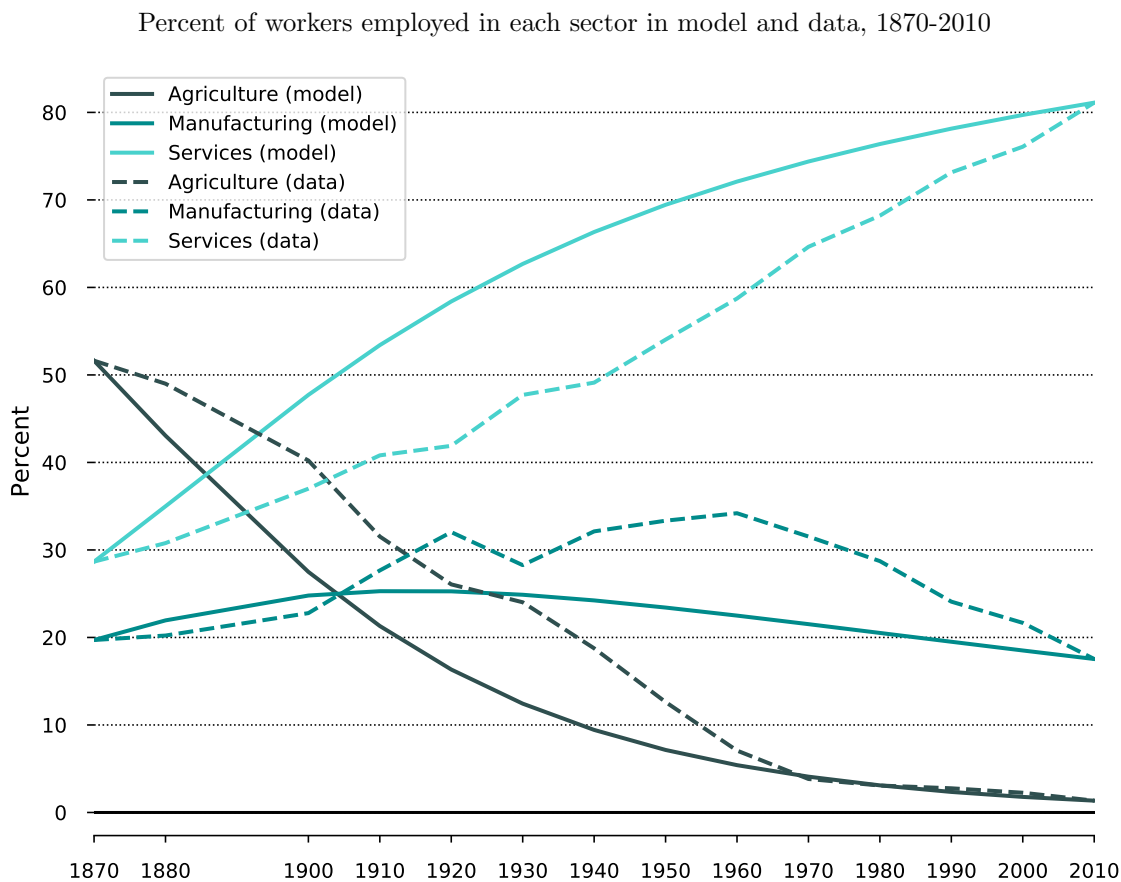
These two normalizations leave six parameters to be pinned down based on historical U.S. data, namely A_m , A_s , g_a , g_m , g_s , and ε . We choose them such that the flexible benchmark model matches the following historical facts for the U.S. economy:

- *Employment shares in 1870 and 2010:* The percent of workers employed in agriculture, manufacturing, and services, at both the beginning, i.e. 1870, and end, i.e. 2010, of the historical sample that we consider. These shares are taken from the Decennial U.S. Censuses digitized and made available by Minnesota Population Center (2017). Because employment shares add up to one, this provides four moment conditions, two in 1870 and two in 2010, to match.
- *Average real GDP per capita growth:* The average annualized growth rate of real GDP per capita from 1870 through 2010, from Bolt and van Zanden (2014). This is the fifth moment we match.
- *Estimates of elasticity of substitution:* Estimates of the elasticity of substitution, ε , based on postwar U.S. national income data reported in Ngai and Pissarides (2008a). This is the final moment condition.

In practice, this means that we follow a procedure similar to the one used by Ngai and Pissarides (2008a). First, just like Ngai and Pissarides (2008a), we choose $\varepsilon = 0.1$ which is consistent with postwar U.S. NIPA data. We then find the values for A_m and A_s under which the employment shares in the flexible benchmark match the 1870 employment shares in the data. Next, we find the relative growth rates of manufacturing and services, i.e. $g_m - g_a$ and $g_s - g_a$, to match the 2010 employment shares. Finally, we choose g_a such that the flexible benchmark matches the average growth rate of real GDP per capita from 1870 through 2010. The values of the non-normalized parameters that we obtain using this method, are $g_a = 0.045$, $g_m = 0.020$, and $g_s = 0.013$ for the annualized TFP growth rates, and $A_m = 2.970$, and $A_s = 1.638$ for the initial relative productivity levels of manufacturing and services.

²This normalization affects the value added shares in GDP, which we do not focus on. It does not affect any of the dynamics that we analyze in this paper.

Figure B.1: Path of employment shares in flexible benchmark and data



Source: Minnesota Population Center (2017) and authors' calculations

In terms of unmatched moments in the data, the implied growth rate of agriculture that results from this calibration is about a percentage point higher than the actual growth rate from merged data from Kendrick (1961) and postwar Industry Accounts of the Bureau of Economic Analysis. When $\varepsilon > 0.1$ the implied growth rate of agriculture is even higher.

Figure B.1 plots the implied time-series paths of the employment shares in the model and the data. These shares are equal in 1870 and 2010 by construction.

B.6 Details about Flexible Benchmark

The flexible benchmark in which (re-)training costs are zero, i.e. $\phi = \gamma_h = 0$ for $h = 1, \dots, H$ is our point of comparison for many of our results. Therefore, we describe the equilibrium path under the flexible benchmark here in more detail. In particular, we focus on the results plotted in Figure B.1. The flexible benchmark equilibrium in this economy is similar to that in Bárány and Siegel (2018) and we touch on many of the same qualitative properties here that Bárány and Siegel (2018)

discuss in much more detail.

Because the outputs of the three sectors are gross complements, the relative wage of workers in the services sector is increasing over time. This is because these workers are getting hired to provide the services that complement that manufacturing and agricultural output for which relatively little labor is need to produce. The trends in wages is depicted in Panel (a) of Row (*i*) of Figure 2.4 in the main text.

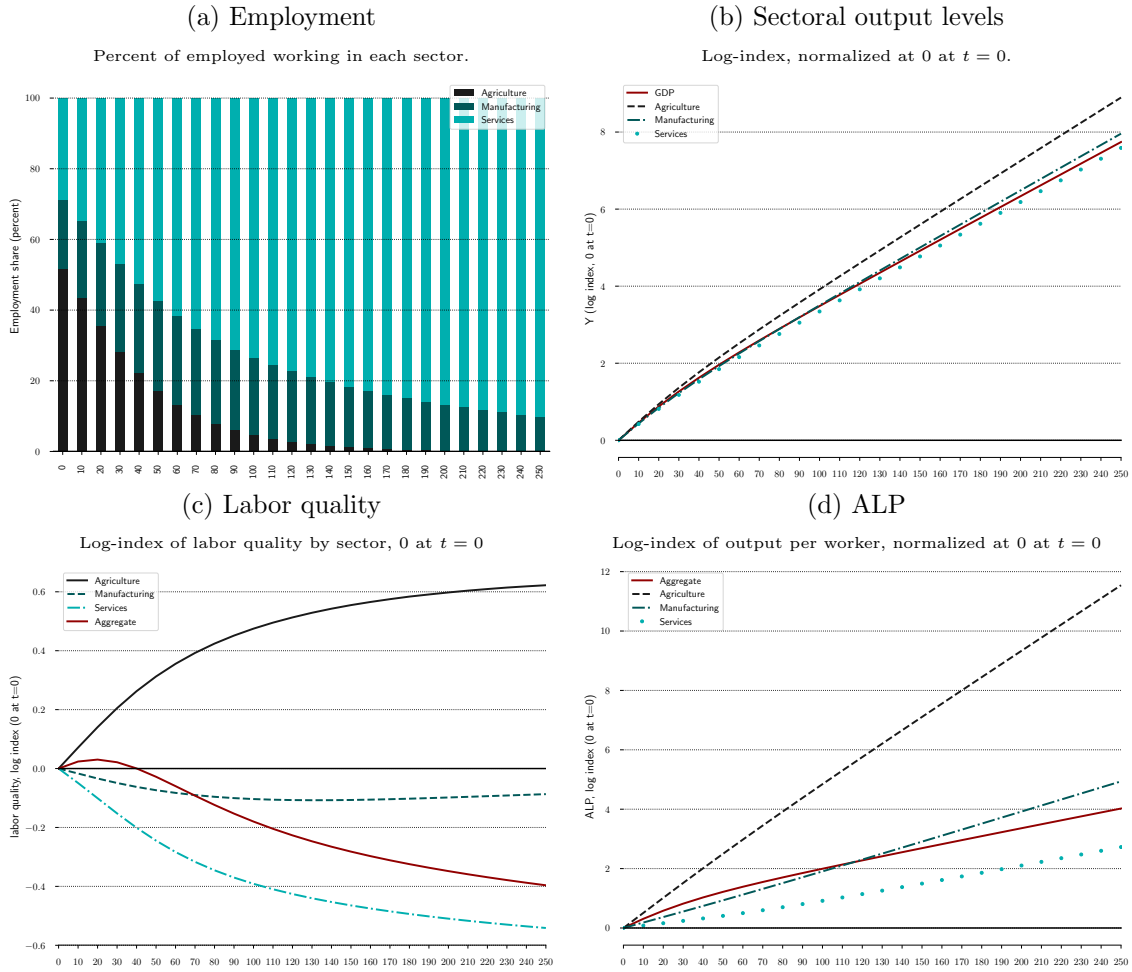
The increasing wage gap draws more and more workers into the service sector over time, as can be seen from Figure B.1a. In the intermediate stages of the transition the employment share of manufacturing peaks while that of agriculture monotonically declines.

Even though a larger fraction of workers is drawn into services over time, this does not translate into higher output growth in services because of the lower TFP growth rate in services, g_s , than in manufacturing, g_m and agriculture, g_a . Moreover, as Ngai and Pissarides (2007a) emphasize, the CES preferences with $\varepsilon < 1$ imply that the relative price, as well as its value-added share in GDP, of the good that grows most slowly will rise over time. As a result, the slowest growing sector makes up an increasing part of GDP over time and drags down overall GDP growth. This can be seen from Figure B.1b where aggregate GDP increasingly aligns with services output over the transition path. In this sense, our model economy suffers from the Cost Disease described in Baumol and Bowen (1968).

But it is not only the exogenously lower rate of TFP growth, g_s , that drags down output growth in the service sector. Average labor productivity growth in this model is determined by both the exogenous rate of TFP growth as well as the selection of workers into different sectors. Just like in Bárány and Siegel (2018) the average quality of workers in the service sector declines over time. This is because the increasing wedge in real wages between services and the other sectors draws workers of lower quality, \tilde{z}_s , into services over time. This can be seen from Figure B.1c, which plots the index of \tilde{z}_i for $i \in \{a, s, m\}$, as well as the average for the whole economy (aggregate). At the end of the quarter millenium transition that we consider this selection of workers across sectors increases labor productivity in agriculture by a factor of two, while reducing that in services by about 40 percent. Because the service sector becomes the predominant sector over time, this decline in labor quality in services drags down aggregate labor quality and, with it, growth in average labor productivity.

This selection effect of workers across sectors is small, however, when compared to the productivity growth rates experienced by each of the sectors. This can be seen by comparing the scale of the vertical axis in Figure B.1c with that of Figure B.1d. The latter plots the log-indices of average labor productivity for the three sectors in the economy as well as aggregate labor productivity. Notice how the Cost Disease of Baumol and Bowen (1968) results in a perpetual productivity slowdown.

Figure B.1: Dynamics of flexible benchmark



Note: Years are plotted on the horizontal axis. $t = 0$ is the equivalent of the beginning of our data sample, i.e. 1870, and $t = 140$ is the equivalent of 2010.

B.7 Implementation of Extended-path Method

For our solution method, we consider the transitional path from $t = 0 - \tilde{t}_l$ until $t = T + \tilde{t}_r$. Here \tilde{t}_l and \tilde{t}_r are padding of the extended path that allow for startup and wind-down periods on the path. We present results for $t = 0, \dots, T$. For $t > T + \tilde{t}_r$ we assume the economy is on a balanced growth path in which $g_i = g > 0$ for $i \in \{a, m, s\}$. Moreover, we assume that, after $t > T + \tilde{t}_r$ workers do not need to spend time on training and retraining anymore. Hence, our solution method isolates the importance of these costs along the transitional path where structural transformation occurs.³

At each point in time, t , the relevant state of the economy consists of three parts. The first are the sector-specific TFP levels, $\{A_{i,t}\}_{i \in \{a,m,s\}}$. The second is the size of the new cohort, $N_t(0)$. The final part is the labor supply, that consists of different age-industry-specific levels $\{\{E_{t-1}(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$.

The TFP levels and initial cohorts evolve exogenously over time, according to

$$A_{i,t} = (1 + g_i) A_{i,t-1}, \text{ where } A_{i,0} \text{ is given and } i \in \{a, m, s\}, \quad (\text{B.35})$$

and

$$N_t(0) = (1 + n) N_{t-1}(0). \quad (\text{B.36})$$

Our solution method is used to solve for the endogenous evolution of $\{\{E_t(i; h)\}_{h=0}^H\}_{i \in \{a,m,s\}}$ in equilibrium along the transitional path for $t = -\tilde{t}_l + 1 \dots T + \tilde{t}_r$. Note that the initial age-industry-specific levels of the labor supply, at time $t = 0$, are given. Thus, we solve the transitional dynamics of the model conditional on the initial state $\{\{E_0(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$. Our method loops over a backward and a forward recursion until reaching convergence.

Backward recursion: Update career choices conditional on labor supply

This recursion starts at time $t = T + \tilde{t}_r$ and runs backwards to $t = -\tilde{t}_l$. It takes the path of the age-industry levels of the labor supply, $\{\{E_t(i; h)\}_{h=0}^H\}_{i \in \{a,m,s\}}$ for $t = -\tilde{t}_l + 1 \dots T + \tilde{t}_r$, as given.

Labor market equilibrium at time t : The main step in this backward recursion is to solve for the equilibrium in the three labor markets at time t , taking as given the age-industry specific levels of the labor supply, $\{\{E_{t-1}(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$, as well as the age-industry specific career continuation values, $\{\{\mathbb{E}_t V_{t+1}(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$. This requires solving for real wages, $\{w_{i,t}\}_{i \in \{a,m,s\}}$, using a system of 3 equations.

Taking as given the continuation values, $\{\{\mathbb{E}_t V_{t+1}(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$, a given set of three wages $\{w_{i,t}\}_{i \in \{a,m,s\}}$ pins down the optimal career decisions that determine $\Phi_t(i)$, $\tilde{\phi}_t(i)$, $\Gamma_t(i, j; h)$, and $\tilde{\gamma}_t(i, j; h)$ through the solution described in Section B.4. With this, we can find the equilibrium wages $\{w_{i,t}\}_{i \in \{a,m,s\}}$ that clear the labor markets, (2.21).

We evaluate $\Phi_t(i)$, $\tilde{\phi}_t(i)$, $\Gamma_t(i, j; h)$, and $\tilde{\gamma}_t(i, j; h)$ at these equilibrium wages and the continuation values $\{\{\mathbb{E}_t V_{t+1}(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$. Then, we iterate backwards.

³We have checked the robustness of our results for the choice of T and g as well as the length of padding \tilde{t}_l and \tilde{t}_r .

Because of the idiosyncratic nature of the (re-)training costs, we can write

$$\mathbb{E}_{t-1}V_t(0) = \sum_{i \in \{a,m,s\}} \Phi_t(i) \left((1 - \phi) \tilde{z}_t(i) w_{i,t} + \frac{1 - \delta}{1 + r} \mathbb{E}_t V_{t+1}(i; 1) \right). \quad (\text{B.37})$$

and for $h = 1, \dots, H - 1$

$$\mathbb{E}_{t-1}V_t(i; h) = \sum_{j \in \{a,m,s\}} \Gamma_t(i, j; h) \left[(1 - \mathbb{I}(i \neq j) \gamma) \tilde{z}_t(i, j; h) w_{i,t} + \frac{1 - \delta}{1 + r} \mathbb{E}_t V_{t+1}(j; h + 1) \right]. \quad (\text{B.38})$$

Finally, we have

$$\mathbb{E}_{t-1}V_t(i; H) = \sum_{j \in \{a,m,s\}} \Gamma_t(i, j; H) (1 - \mathbb{I}(i \neq j) \gamma) \tilde{z}_t(i, j; H) w_{j,t}. \quad (\text{B.39})$$

These equations now allow us to solve the relevant continuation values using a backward recursion. Of course, this solution is conditional on $\{\{E_t(i; h)\}_{h=0}^H\}_{i \in \{a,m,s\}}$. With these continuation values for $t - 1$ in hand we can now solve the labor market equilibrium for period $t - 1$ conditional on $\{\{E_{t-1}(i; h)\}_{h=0}^H\}_{i \in \{a,m,s\}}$ using the same method and role this backward recursion all the way from $t = T + \tilde{t}_r$ to $t = -\tilde{t}_l$.

Forward recursion: Update labor supply conditional on career choices

In the forward recursion we now update the path of the age-industry specific levels of the labor supply, $\{\{E_t(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$, using the career choices solved in the backward recursion. This simply involves iterating over the law of motion of the labor supply from (2.17) and (2.18). This recursion is initialized using the initial condition that gives the age-industry specific labor supply levels at time $t = 0$, i.e. $\{\{E_0(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$. We continue to loop over the backward and forward recursions until the calculated path of the age-industry specific labor supply levels, $\{\{E_t(i; h)\}_{h=1}^H\}_{i \in \{a,m,s\}}$ for $t = -\tilde{t}_l + 1 \dots T + \tilde{t}_r$, converges.

APPENDIX C
APPENDICES TO CHAPTER 3

C.1 Data Sources

The data I use comes from three sources: the ILOSTAT database from the International Labor Organization, the Penn World Tables, and the census databases from the IPUMS-International Project. The macroeconomic variables, such as employment and GDP, come from the PWT. The data for the distribution of employment shares across occupations is a combination of ILOSTAT and consolidated data from IPUMS-International. Both IPUMS and ILOSTAT have occupational information at the ISCO-08 1-digit level, which I combine as follows. Abstract occupations consist of Managers, Professionals, and Technicians and Associate Professionals (groups 1, 2, and 3). Routine occupations consist of Clerical Support Workers, Craft and Related Trades Workers, and Plant and Machine Operators, and Assemblers (groups 4, 7, and 8). Manual occupations consist of Services and Sales Workers, Skilled Agricultural, Forestry, and Fishery Workers, and Elementary Occupations (groups 5, 6, and 9). When data from IPUMS and ILOSTAT overlay, the data from IPUMS takes priority, and the rest of the series is built with the changes in the employment shares.

Table C.1: Country Data Availability

Country	Years	Country	Years
Argentina	2004–2014	Armenia	2011–2014
Aruba	1994–2011	Australia	1991–2014
Austria	1981–2014	Azerbaijan	2003–2014
Bahamas	1991–2009	Bahrain	1991–2004
Barbados	1994–2014	Belarus	1999–2009
Belgium	1993–2014	Belize	1993–1999
Bermuda	2000–2007	Bhutan	2006–2014
Bolivia	1992–2014	Botswana	1991–2011
Brazil	1980–2014	Bulgaria	2000–2014
Burkina Faso	1985–1996	Cambodia	1998–2008
Canada	1981–2011	Cayman Islands	1991–2008
Chile	1982–2014	China	1982–1990
Hong Kong SAR	1994–2014	Costa Rica	1984–2014
Croatia	1996–2014	Cyprus	1999–2014
Czech Republic	1993–2014	Denmark	1992–2014
Dominica	1991–2001	Dominican Republic	1981–2014
Ecuador	1982–2014	Egypt	1986–2014
El Salvador	2008–2012	Estonia	1990–2014
Fiji	1986–2007	Finland	1997–2014
France	1982–2014	Georgia	1998–2007
Germany	1992–2014	Ghana	1984–2010
Greece	1981–2014	Guatemala	2012–2014
Guinea	1983–1996	Haiti	1982–2003
Hungary	1990–2014	Iceland	1991–2014
India	1983–2012	Indonesia	1980–1995
Iran	1996–2014	Ireland	1981–2014
Israel	1995–2014	Italy	1992–2014
Jamaica	1982–2008	Japan	2009–2014
Kazakhstan	2001–2013	Kyrgyzstan	2003–2014
Latvia	1996–2014	Lebanon	2004–2007
Lithuania	1997–2014	Luxembourg	1992–2014

Table C.2: Country Data Availability (Ctn'd)

Country	Years	Country	Years
Malawi	1987–2008	Malaysia	1980–2014
Mali	1987–2014	Malta	2000–2014
Mauritius	1995–2014	Mexico	1990–2014
Mongolia	2005–2014	Montenegro	2005–2014
Morocco	1982–2011	Mozambique	1997–2007
Namibia	2000–2014	Nepal	1999–2008
Netherlands	1992–2014	New Zealand	1992–2008
Nicaragua	1995–2014	Norway	1996–2014
Pakistan	2002–2014	Panama	1980–2014
Paraguay	1982–2014	Peru	1993–2014
Philippines	2001–2014	Poland	1995–2014
Portugal	1991–2014	Republic of Moldova	1999–2014
Romania	1995–2014	Russian Federation	1997–2014
Saint Lucia	1994–2006	Sao Tome and Principe	2003–2012
Senegal	1988–2002	Serbia	2004–2014
Seychelles	2011–2014	Singapore	1985–2014
Slovakia	1994–2014	Slovenia	1995–2014
South Africa	1996–2014	Spain	1981–2014
Sri Lanka	2002–2014	Suriname	2004–2014
Sweden	1997–2014	Switzerland	1980–2014
Macedonia	2002–2014	Taiwan	1994–2013
Thailand	2001–2013	Trinidad and Tobago	1980–2014
Turkey	1985–2014	Turks and Caicos Islands	2002–2007
Tanzania	1988–2012	Uganda	1991–2003
Ukraine	1999–2014	United Kingdom	1991–2014
United States	1980–2014	Uruguay	1996–2014
Venezuela	1981–2001	Vietnam	1999–2014
Yemen	1999–2014	Zambia	2000–2010
Zimbabwe	2011–2014		

C.2 Shift-share Details

This section documents how world polarization happens mostly within industries. We can express the share of each occupation as a weighted average at the country level. For period t :

$$l_{c,j,t} = \sum_I s_{c,t}(I) l_{c,j,t}(I) \quad (\text{C.1})$$

where $l_{c,j,t}$ is the country's employment share of occupation j , $s_{c,t}(I)$ is the country's employment share of industry I , and $l_{c,j,t}(I)$ is the share of occupation j within industry I . The change between period 0 and t can be decomposed into its *between* and *within* industry components:

$$\Delta l_{c,j,t} = \underbrace{\sum_I \Delta s_{c,t}(I) \bar{l}_{c,j,t}(I)}_{\text{Between industries effect}} + \underbrace{\sum_I \Delta l_{c,j,t}(I) \bar{s}_{c,t}(I)}_{\text{Within industries effect}} \quad (\text{C.2})$$

$\bar{l}_{c,j,t}(I)$ is the average between time 0 and t of the conditional occupation share, and $\bar{s}_{c,t}(I)$ that of the industry share. The *between* effect refers to the impact of structural transformation, the changes in the productive structure of the economy. The *within* effect refers to the occupational mix inside each industry.

Both the changes within- and between-industries contribute to polarization. The within-industry changes dominate: on average, the within-industry contribution accounts for 63% of polarization.

C.3 VAR Convergence System

The VAR model analyzes the productivity *gaps* with respect to the U.S. levels. Remember these gaps are defined as:

$$\tilde{A}_{c,j,t} = \frac{A_{c,j,t}}{A_{US,j,t}}$$

The vector of log deviations from the U.S. is $\tilde{\mathbf{a}}_{c,t} = [\log \tilde{A}_{a,c,t}, \log \tilde{A}_{r,c,t}, \log \tilde{A}_{m,c,t}]'$, and the econometric model I use is

$$\tilde{\mathbf{a}}_{c,t} - \tilde{\mathbf{a}}_{c,t-1} = \boldsymbol{\alpha}_c + \mathbf{B} \tilde{\mathbf{a}}_{c,t-1} + \mathbf{e}_{c,t} \quad (\text{C.3})$$

$\boldsymbol{\alpha}_c = [\alpha_{a,c}, \alpha_{r,c}, \alpha_{m,c}]$ captures country fixed-effects, \mathbf{B} is a 3×3 matrix containing the β coefficients associated with the lagged log-deviations, and $\mathbf{e}_{c,t} = [\varepsilon_{a,c,t}, \varepsilon_{r,c,t}, \varepsilon_{m,c,t}]$ is the vector of error terms.

This system implies a long-run expectations for the gaps in each of the countries

$$\tilde{\mathbf{a}}_{c,t} = \boldsymbol{\alpha}_c + (\mathbf{I} + \mathbf{B}) \tilde{\mathbf{a}}_{c,t-1} \quad (\text{C.4})$$

The dynamics are dictated by the eigenvalues associated to $\mathbf{I} + \mathbf{B}$. When all eigenvalues are unique, the system can be transformed into three autonomous difference equations. Matrix $\mathbf{I} + \mathbf{B}$ can be diagonalized as follows:

$$\mathbf{D} = \mathbf{V}^{-1}(\mathbf{I} + \mathbf{B})\mathbf{V} \quad (\text{C.5})$$

where \mathbf{D} is a diagonal matrix containing the eigenvalues, and \mathbf{V} is an invertible matrix with the associated eigenvectors. If we define, $\mathbf{z}_{c,t} = \mathbf{V}^{-1}\tilde{\mathbf{a}}_{c,t}$, the system can be rewritten as:

$$\begin{aligned}\mathbf{z}_{c,t} &= \mathbf{V}^{-1}\tilde{\mathbf{a}}_{c,t} \\ &= \mathbf{V}^{-1}\boldsymbol{\alpha}_c + \mathbf{V}^{-1}(\mathbf{I}+\mathbf{B})\mathbf{V}\mathbf{V}^{-1}\tilde{\mathbf{a}}_{c,t-1} \\ &= \boldsymbol{\zeta}_c + \mathbf{D}\mathbf{z}_{c,t-1}\end{aligned}\tag{C.6}$$

This is a system of autonomous difference equations. If the j -th row has a real-valued eigenvalue (λ_j), then

$$\begin{aligned}z_{j,c,t} &= \zeta_{j,c} + \lambda_j z_{j,c,t-1} \\ \Rightarrow z_{j,c,t} &= \frac{\zeta_{j,c}}{1 - \lambda_j} + \left(z_{j,0} - \frac{\zeta_{j,c}}{1 - \lambda_j} \right) \lambda_j^t\end{aligned}\tag{C.7}$$

This equation is convergent when $|\lambda_j| < 1$, so that the last term goes to 0 as $t \rightarrow \infty$. $\zeta_{j,c}/(1 - \lambda_j)$ is the steady-state level, and λ_j determines how fast deviations from it converge. Therefore, $1 - \lambda_j$ is the percentage of the gap with respect to the steady-state level that is closed each period, or the convergence rate.

If the j -th and $j + 1$ -th rows have complex-valued eigenvalues of the form $\lambda_j = \alpha + \theta i$, $\lambda_{j+1} = \alpha - \theta i$, then the homogeneous parts have form

$$z_{j,c,t} = r^t(A_1 \cos \omega t + A_2 \sin \omega t)\tag{C.8}$$

$$z_{j+1,c,t} = r^t(B_1 \cos \omega t + B_2 \sin \omega t)\tag{C.9}$$

where $r^2 = \alpha^2 + \theta^2$ is the modulus, and defines the convergence rate.

C.4 A Model of Task Investment Specific Technical Change to Understand Polarization

The goal of this model is to understand two aspects of modern economic growth: the technological factors behind occupational employment shares, and their change over time. This fits naturally into an investment specific technical change framework, which needs to be extended to include the different tasks.

Economic Environment

Consider an economy populated by a representative household. Its preferences are given by:

$$U = \sum_{t=0}^{\infty} \beta^t \log(c_t)\tag{C.10}$$

where c_t represents consumption during period t , and $\beta > 0$ is the discount factor.

This economy follows the task based approach to production (Acemoglu and Autor, 2011). Tasks are the base units of work, that are combined as intermediate inputs

to produce final output. There are three tasks: abstract (a), routine (r), and manual (m). Task specific variables are denoted by the j subindex, so that $j \in \{a, r, m\}$. The production of task j requires the labor services of occupation j and capital, which is specifically tailored to the production of that task. Technological progress is captured by A_t , which is total factor productivity, and is neutral across tasks. These are combined according to a Cobb-Douglas production function:

$$y_{j,t} = A_t k_{j,t}^\alpha l_{j,t}^{1-\alpha} \quad (\text{C.11})$$

where $\alpha > 0$ determines the capital share in income.

Final output requires three intermediate inputs: abstract, routine, and manual tasks. These are combined according to the following constant elasticity of substitution production function, similar to Goos *et al.* (2014):

$$y_t = \left(\sum_{j \in \{a, r, m\}} \omega_j^\sigma y_{j,t}^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}} \quad (\text{C.12})$$

where $\varepsilon > 0$ is the elasticity of substitution among tasks, $\omega_j > 0$ is the production intensity of task j , and $\sigma = (1 - \alpha + \alpha\varepsilon)/\varepsilon$ is a scaling factor due to the nested productive structure. These task intensities add up to one.

Final output can be split between consumption and investment in the three capital stocks:

$$y_t = c_t + \sum_{j \in \{a, r, m\}} i_{j,t} \quad (\text{C.13})$$

Investment expenditures do not translate one-to-one into new capital. Instead, one unit of investment converts into q_j units of capital. The capital accumulation evolves according to:

$$k_{j,t+1} = (1 - \delta)k_{j,t} + q_{j,t}i_{j,t} \quad (\text{C.14})$$

where δ is the depreciation rate, and $q_{j,t}$ is the task investment specific technological level. This notion of investment specific technical change follows Greenwood *et al.* (1997). Differently to them, investment specific technical change is *task specific*, which means that the efficiency to produce capital is different across tasks. This also implies that capital is task specific; it can only be used in the production of task j . Capital stocks are convertible across tasks, but the household faces different prices when doing so ($1/q_{j,t}$ units of consumption).

Labor can be allocated to the production of the three tasks. It is also perfectly mobile among tasks, and homogeneous. The total labor force is normalized to 1, so that:

$$1 = \sum_{j \in \{a, r, m\}} l_{j,t} \quad (\text{C.15})$$

Competitive Equilibrium

In each period, the state of the economy is characterized by total factor productivity A_t , and by the task specific capital stocks $\{k_{j,t}\}_{j \in \{a,r,m\}}$ and investment specific technological levels $\{q_{j,t}\}_{j \in \{a,r,m\}}$. Their sequence of technological levels determines the sequential competitive equilibrium, which is explained in what follows.

Household

The household owns the task specific stocks of capital, which are rented at the price $R_{j,t}$. It also supplies its labor to the task producers, and earns wages $w_{j,t}$. The problem it faces is:

$$\begin{aligned} & \max_{\{c_t, \{l_{j,t}, i_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \beta^t \log(c_t) \\ & \text{subject to:} \\ & c_t + \sum_{j \in \{a,r,m\}} i_{j,t} = \sum_{j \in \{a,r,m\}} (w_{j,t} l_{j,t} + R_{j,t} k_{j,t}) \\ & k_{j,t+1} = (1 - \delta) k_{j,t} + q_{j,t} i_{j,t} \quad \text{for } j \in \{a, r, m\} \\ & 1 = \sum_{j \in \{a,r,m\}} l_{j,t} \end{aligned} \quad (\text{C.16})$$

taking the sequence of prices $\{\{R_{j,t}, w_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^{\infty}$ as given. Consumption goods work as the numeraire in this model.

Intermediate Task Producers

Firms producing tasks sell their product and rent capital and labor in competitive markets. Their optimization problem is:

$$\max_{k_{j,t}, l_{j,t}} \pi_{j,t} = p_{j,t} A_t k_{j,t}^{\alpha} l_{j,t}^{1-\alpha} - R_{j,t} k_{j,t} - w_{j,t} l_{j,t} \quad (\text{C.17})$$

taking prices $p_{j,t}$, $R_{j,t}$ and $w_{j,t}$ as given. This is a static problem, and due to constant returns to scale and competitive markets, profits are zero.

Final Output Producers

Firms producing final output also participate in competitive markets. Their optimization problem is:

$$\max_{\{y_{j,t}\}_{j \in \{a,r,m\}}} \pi_t = \left(\sum_{j \in \{a,r,m\}} \omega_j^{\sigma} y_{j,t}^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} - \sum_{j \in \{a,r,m\}} p_{j,t} y_{j,t} \quad (\text{C.18})$$

taking task prices $p_{j,t}$ as given. As in the problem of intermediate task producers, this is a static problem where profits are zero.

Equilibrium

A sequential competitive equilibrium in this economy is a sequence of allocations $\{c_t, y_t, \{l_{j,t}, i_{j,t}, k_{j,t+1}, y_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^{\infty}$, a sequence of prices $\{\{R_{j,t}, w_{j,t}, p_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^{\infty}$, a sequence of task investment specific technological levels $\{\{q_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^{\infty}$, and a sequence of total factor productivities $\{A_t\}_{t=0}^{\infty}$ such that:

1. The sequence $\{c_t, \{l_{j,t}, i_{j,t}, k_{j,t+1}\}_{j \in \{a,r,m\}}\}_{t=0}^\infty$ solves the household's optimization problem, taking as given the initial stocks of capital $\{k_{j,0}\}_{j \in \{a,r,m\}}$, and the sequence of prices and investment technological levels $\{\{R_{j,t}, w_{j,t}, q_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^\infty$.
2. The sequence $\{\{y_{j,t}, k_{j,t}, l_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^\infty$ solves the problems of intermediate task producers, taking as given the sequence of prices and total factor productivities $\{\{R_{j,t}, w_{j,t}, p_{j,t}\}_{j \in \{a,r,m\}}, A_t\}_{t=0}^\infty$.
3. The sequence $\{y_t, \{y_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^\infty$ solves the problems of final output firms, taking as given the sequence of prices $\{\{p_{j,t}\}_{j \in \{a,r,m\}}\}_{t=0}^\infty$.
4. Markets clear:

$$y_t = c_t + \sum_{j \in \{a,r,m\}} i_{j,t} \quad (\text{C.19})$$

$$y_{j,t} = A_t k_{j,t}^\alpha l_{j,t}^{1-\alpha} \quad (\text{C.20})$$

$$1 = \sum_{j \in \{a,r,m\}} l_{j,t} \quad (\text{C.21})$$

5. Capital stocks follow their laws of motion

$$k_{j,t+1} = (1 - \delta)k_{j,t} + q_{j,t}i_{j,t} \quad (\text{C.22})$$

C.4.1 Model Estimation

This section describes how the model is parametrized to match certain features of the data. This is done in several steps. Broadly speaking, the employment shares are informative of the relative composition of task specific capital stocks. These are leveled up using data on real income and total factor productivity. The capital accumulation process implies a path for the units of investment, that are a mixture of the resources invested, and the levels of investment specific technology $(i_{j,t}, q_{j,t})$. To separate those, I use data on investment shares for a benchmark year. Finally, the intertemporal Euler equations imply the rest of the path for the levels of task investment specific technology.

At each point in time, occupational employment shares in the data determine the model's relative capital stocks. This follows from the optimality conditions when task producers decide how much labor to hire. Labor is freely mobile and homogeneous, which equalizes wages across the production of tasks. Therefore, the first order conditions for the task producers' problem (C.17) imply that

$$\begin{aligned} \frac{l_{i,t}}{l_{j,t}} &= \left(\frac{p_{i,t}}{p_{j,t}} \right)^{\frac{1}{\alpha}} \frac{k_{i,t}}{k_{j,t}} \\ &= \frac{\omega_i}{\omega_j} \left(\frac{k_{j,t}}{k_{i,t}} \right)^{\frac{\alpha(1-\varepsilon)}{1-\alpha(1-\varepsilon)}} \end{aligned} \quad (\text{C.23})$$

With information on the production intensities, ω_j , and the capital share α and elasticity of substitution ε , this equation leads to the relative capital stocks.

The levels of capital are used to match the levels of income in the data. Solving for the intertemporal allocations of the model in terms of the capital stocks yields the following expression for final output:

$$y_t = A_t \left(\sum_{j \in \{a,r,m\}} \omega_j / k_{j,t}^{\frac{\alpha(1-\varepsilon)}{1-\alpha(1-\varepsilon)}} \right)^{-\frac{1-\alpha(1-\varepsilon)}{1-\varepsilon}} \quad (C.24)$$

The income and total factor productivity data, that come from the Penn World Tables, are normalized to the US levels in 2000. The next logical step is to normalize the US capital stocks in that year. This serves two purposes. First, it provides a way to estimate the intensity parameters ω_j from equation (C.23). Second, this choice expresses the rest of the capital stocks in relative terms to US levels in 2000. These are recovered from equation (C.24).

Capital accumulation provides a way to measure effective investment:

$$k_{j,t+1} = (1 - \delta)k_{j,t} + q_{j,t}i_{j,t} \quad (C.25)$$

The problem now is to disentangle investment technological levels and investment expenses from effective investment, i.e., separating the $q_{j,t}i_{j,t}$ series into $q_{j,t}$ and $i_{j,t}$. As a first approximation, I use the Euler equations:

$$\frac{c_{t+1}}{c_t} = q_{j,t}\beta \left[R_{j,t+1} + \frac{1 - \delta}{q_{j,t+1}} \right] \quad (C.26)$$

Assuming that the growth rate of $q_{j,t}$ is the same for all tasks for a year provides a way to approximate the task investment specific technological levels:

$$q_{j,t}R_{j,t+1} = q_{j',t}R_{j',t+1} \quad (C.27)$$

Finally, the levels are scaled to match the investment shares in 2000.

This solves for the task investment specific technology levels in one year. The path is completed through an iterative process that uses the household's Euler equations and the resource constraints. The Euler equations inform about relative growth rates in $q_{j,t}$, and the resource constraints level that growth.

The initial estimates of $q_{j,t}$ imply as well the consumption level. Therefore, the Euler equations and the resource constraint for the following period pose a system of 4 equations in 4 unknowns:

$$\frac{c_{t+1}}{c_t} = q_{j,t}\beta \left[R_{j,t+1} + \frac{1 - \delta}{q_{j,t+1}} \right] \quad (C.28)$$

$$c_{t+1} = y_{t+1} - \sum_{j \in \{a,r,m\}} \frac{q_{j,t+1}i_{j,t+1}}{q_{j,t+1}} \quad (C.29)$$

This is possible because the series of capital stocks was already estimated: it determines the rental rates $R_{j,t+1}$, output y_{t+1} , and effective investment $i_{j,t+1}q_{j,t+1}$. This

process can be iterated forwards, up to the final year with information, or backwards to the initial year. This way, the entire path of task investment specific technology levels and consumption can be inferred.

The rest of the parameters are determined from the Penn World Tables, or borrowed from other studies. The capital share in income is set to 0.4, which is the average found in the data. Similarly, the depreciation rate is averaged to 0.05. The discount factor β is set to 0.95, and the elasticity of substitution is set to 0.0625, to match the estimate of Vindas (2017) in the United States.

C.4.2 Alternative Identification Strategy

An alternative way to disentangle initial task investment specific technological levels is explained in this section. Its actual implementation is left to future iterations of this project.

Investment is measured in different units, depending on the context. First, the investment expenses $i_{j,t}$, like those on the resource constraint (C.13), are denoted in units of *consumption*. Second, effective investment is measured in units of capital, which is $q_{j,t}i_{j,t}$. The relative price of investment for capital in task j is the number of units of consumption paid, per measured units of investment:

$$\begin{aligned}\frac{P_{j,t}^I}{P_{j,t}^C} &= \frac{i_{j,t}}{q_{j,t}i_{j,t}} \\ &= \frac{1}{q_{j,t}}\end{aligned}\tag{C.30}$$

In the data, we don't observe these individually, but rather through an aggregate investment price. This is a weighted average over investment expenditure shares, which is the first moment to target. It is given by:

$$\begin{aligned}\frac{P_t^I}{P_t^C} &= \sum_{j \in \{a,r,m\}} \frac{i_{j,t}}{\sum_{j' \in \{a,r,m\}} i_{j',t}} \frac{P_{j,t}^I}{P_{j,t}^C} \\ &= \sum_{j \in \{a,r,m\}} \frac{i_{j,t}}{\sum_{j' \in \{a,r,m\}} i_{j',t}} \frac{1}{q_{j,t}}\end{aligned}\tag{C.31}$$

The second moment is the aggregate relative price of capital. It is informative since this is another weighed average; the composition of the current capital stock is different from the composition of investment expenses. In particular, this is equal to the replacement-value weighted average of the relative prices of capital, and is given by:

$$\frac{P_t^K}{P_t^C} = \sum_{j \in \{a,r,m\}} \frac{k_{j,t}/q_{j,t}}{\sum_{j' \in \{a,r,m\}} k_{j',t}/q_{j',t}} \frac{A_{t+1}^{\frac{1}{\alpha}}}{q_{j,t}}\tag{C.32}$$

The third and final moment is the investment rate, which is given by:

$$s_t = \frac{\sum_{j \in \{a,r,m\}} i_{j,t}}{Y_t} \quad (\text{C.33})$$

Intuitively, the investment and capital prices pin down the relative levels of $q_{j,t}$ across occupations, since these are two different weighted averages. The investment rate, on the other hand, defines the actual levels that are compatible with the aggregate investment decisions.

To solve this system, it has to be expressed in terms of the quantities that are available so far. This requires substituting $i_{j,t}$ with $q_{j,t}i_{j,t}/q_{j,t}$. This results in a quadratic set of equations in the task investment specific technology levels. To solve it, I log-linearize this system around 1. Its approximation is given by:

$$\ln \left(\frac{P_t^I}{P_t^C} \right) \approx - \sum_j \frac{i_{j,t}q_{j,t}}{\sum_{j'} i_{j',t}q_{j',t}} \left[2 - \frac{i_{j,t}q_{j,t}}{\sum_{j'} i_{j',t}q_{j',t}} \right] \ln q_{j,t} \quad (\text{C.34})$$

$$\ln \left(\frac{P_t^K}{P_t^C} \right) \approx - \sum_j \frac{k_{j,t}}{\sum_{j'} k_{j',t}} \left[2 - \frac{k_{j,t}}{\sum_{j'} k_{j',t}} \right] \ln q_{j,t} \quad (\text{C.35})$$

$$s_t \approx \sum_j \frac{i_{j,t}q_{j,t}}{Y_t} - \sum_j \frac{i_{j,t}q_{j,t}}{Y_t} \ln q_{j,t} \quad (\text{C.36})$$

This is a linear system on log-deviations from 1, which then allows to approximate the task investment specific technological levels in 2005.