Job Polarization and Structural Change*

Zsófia L. Bárány† Christian Siegel‡

26 November 2014

Abstract

Job polarization is a widely documented phenomenon in developed countries since the 1980s: employment has been shifting from middle to low- and high-income workers, while average wage growth has been slower for middle-income workers than at both extremes. We document that polarization has started as early as the 1950s in the US, and that this process is closely linked to the shift from manufacturing to services. Based on these observations we propose a structural change driven explanation for polarization. In order to analyze not only the evolution of employment shares, but also of relative wages, we extend one of the standard frameworks of structural change by modeling the sectoral choice of workers in a Roy-type setup. Introducing this novel feature does remarkably well in matching the sectoral labor market outcomes over the past 60 years.

JEL codes: E24, J22, O41

Keywords: Job Polarization, Structural Change, Roy model

*We wish to thank Francesco Caselli, David Dorn, Georg Duernecker, Georg Graetz, Per Krusell, Silvana Tenreyro and Akos Valentinyi as well as seminar and conference participants at Trinity College Dublin, UC Berkeley, Sciences Po, Stockholm School of Economics, Exeter, the Vienna Joint Macro Seminar, the 7th Nordic Summer Symposium in Macroeconomics, the European Meeting of the Econometric Society in Gothenburg, the 16th IZA/CEPR European Summer Symposium in Labour Economics, the 2013 Cologne Workshop on Macroeconomics, the 2013 Winter Meeting of the Hungarian Society of Economists, the Christmas Meetings of the German Economists Abroad in Konstanz, and the German Economic Association Annual Meeting 2014 for useful comments. Latest version available at: `https://dl.dropboxusercontent.com/u/33862610/paper_jobpol.pdf?dl=1`

†Sciences Po. Email: zsofia.barany@sciencespo.fr

‡University of Exeter. Email: c.siegel@exeter.ac.uk

1 Introduction

The polarization of the labor market in terms of occupations is a widely documented phenomenon in the US and several European countries since the 1980s.[1] This phenomenon, besides the relative growth of wages and employment of high earning occupations, also entails the relative growth of wages and employment of low earning occupations. The leading explanation for polarization is the routinization hypothesis, which relies on the assumption that information and computer technologies (ICT) substitute for middle-skill and hence middle-earning (routine) occupations, whereas they complement the high-skilled and high-earning (abstract) occupations (Autor, Levy, and Murnane (2003), Autor, Katz, and Kearney (2006), Autor and Dorn (2013), Michaels, Natraj, and Van Reenen (2014), Goos, Manning, and Salomons (2014)).

The contribution of our paper is twofold. First, we document a set of facts which raise flags that ICT and globalization, although certainly playing a role from the 1980s onwards, might not be the only driving forces behind this phenomenon. Second, based on these facts we propose a novel perspective on the polarization of the labor market, one based on structural change.

Our analysis of US Census data for the period 1950-2000 and American Community Survey (ACS) data for 2007 reveals some novel facts.[2] First, while most of the literature on polarization focuses on occupations, we document that labor market polarization is present also in terms of broadly defined sectors: low-skilled and high-skilled service workers, who are at opposite ends of the earnings distribution have been gaining in terms of wages and employment at the expense of manufacturing workers. Second, sectoral shifts play a significant role in the observed occupational employment changes. Third, we find that polarization, both in terms of occupations and sectors, has started as early as the 1950-1960s in the US. This implies that polarization started long before ICT or increased trade flows could have impacted the labor market. Observing a) that polarization seems to be a long-run phenomenon, b) the contraction in manufacturing seems to be contributing to the decline in routinizable jobs, c) that manufacturing employment started to fall, while service employment started to increase in the 1950s-1960s, it is natural to investigate whether the structural shift of the economy is driving the polarization of the labor market.

Based on these observations we propose a structural change driven explanation for the joint polarization of wages and employment. We introduce a Roy-type selection mechanism into a multi-sector growth model, where each sector values a specific skill. Individuals, who are heterogeneous along a range of skills, optimally select which sector to work in. As long as the goods produced by the different sectors are not perfect substitutes, a change in relative productivities triggers a change in labor demand across sectors. If the elasticity of substitution is less than one, then labor demand increases in the relatively slow growing sectors, otherwise in the relatively fast growing sectors. The remuneration in the

[1]See for example Autor, Katz, and Kearney (2006) and Autor and Dorn (2013) for the US, Goos and Manning (2007) for the UK, Spitz-Oener (2006) and Dustmann, Ludsteck, and Schönberg (2009) for Germany, Goos, Manning, and Salomons (2009) for Europe.

[2]Analyzing the data until more recent years does not affect our findings; we chose 2007 as the final year to exclude the potential impact of the financial crisis.

sectors with expanding labor demand has to increase in order to attract more workers.

In particular we assume that there are three types of consumption goods: low-end service, manufacturing and high-end service goods. Splitting services into two categories is guided by considerations both on the production and the consumption side. Defining low-end services as those that could be home produced, we see in the data that workers in low- and high-end services have very different characteristics and earn very different wages. Guided by this we assume that low-end services require raw labor in production (for this reason we refer to them as low-skilled services), while high-end services and manufacturing make use of sector specific skills.[3] From a consumption perspective we believe that low- and high-end services are just as substitutable with each other as they are with manufacturing goods, rather than being perfect substitutes, as would be implied by a single service sector.

We show that if goods and the two types of services are complements, then as relative labor productivity in manufacturing increases, labor has to reallocate from manufacturing to both service sectors in order to meet the demand for high- and low-skilled services. To attract more workers into these two sectors, their wages have to improve relative to manufacturing. Given that manufacturing jobs tend to be in the middle of the wage distribution, this mechanism leads to a pattern of polarization.

We calibrate the model to quantitatively assess the contribution of structural change – driven by unbalanced technological progress – to the polarization of wages and employment. Taking labor productivity growth from the data and using existing estimates for the elasticity of substitution between sectors, we find that our model predicts around two thirds of the relative average wage gain of high- and low-skilled services compared to manufacturing, and our predictions for the employment shares are within a 20 per cent range of their values in the data.

This paper builds on and contributes to the literature both on polarization and on structural change. To our knowledge, these two phenomena until now have been studied separately.[4] However, according to our analysis of the data, polarization of the labor market and structural change are closely linked to each other, and according to our model, industrial shifts can lead to polarization.

The structural change literature has documented for several countries that as income increases employment shifts away from agriculture and from manufacturing towards services, while expenditure shares follow similar patterns (Kuznets (1957), Maddison (1980), Herrendorf, Rogerson, and Valentinyi (2014)). In particular the employment and expenditure share of manufacturing has been declining since the 1950s in the US, while those in services have been increasing. From an empirical perspective, we add to this literature by documenting that in the US since the 1950s the employment patterns are mimicked by the path of relative average wages. The economic forces that lead to structural transformation are related to either preferences or technology. The preference explanation relies on non-homothetic preferences, such that changes in aggregate income lead to a reallocation of employment across sec-

[3] Buera and Kaboski (2012) also split services into low- and high-skilled: their selection is based on the fraction of college educated workers in the industry. Their main interest is linking the rising skill premium to the increasing share of services in value added, and they emphasize the home vs market production margin. Our focus is very different: sectoral wages.

[4] Some papers on polarization look at the contribution of between industry shifts to the polarization of occupational employment, but not at what causes these between industry shifts.

tors (Kongsamut, Rebelo, and Xie (2001), Boppart (2014)). The mechanisms related to technology rely either on differential total factor productivity (TFP) growth across sectors (Ngai and Pissarides (2007)) or on changes in the supply of an input used by different sectors with different intensities (Caselli and Coleman (2001), Acemoglu and Guerrieri (2008)).

We build on the model of Ngai and Pissarides (2007) closely, with one important modification: we explicitly model sectoral labor supply. As our goal is to study the joint evolution of employment and wages, we introduce heterogeneity in workers' skills, who endogenously sort into different sectors. In order to meet increasing labor demands in certain sectors – driven by structural change – the relative wages of those sectors have to increase. Since we model the sector of work choice, we can analyze the effects of structural change on relative sectoral wages, which is not common in models of structural change.[5] Another modification of Ngai and Pissarides (2007) is that we do not model capital, as our interest is in the heterogeneity of labor supply. This also implies that a change in relative sectoral labor productivity can be driven by differential sectoral TFP changes or by capital accumulation and different sectoral capital intensities.[6]

The polarization literature typically focuses on employment and wage patterns after the 1980s or 1990s. We contribute to this literature by documenting that in the US the polarization of occupations in terms of wages and employment has started as early as the 1950s. One explanation suggested in the literature are consumption spillovers. This argument suggests that as the income of high-earners increases, their demand for low-skilled service jobs increases as well, leading to a spillover to the lower end of the wage distribution (Manning (2004), Mazzolari and Ragusa (2013)). The other leading explanation is routinization.[7] The routinization hypothesis suggests, that it is ICT that leads to the substitution of routine (middle-earning) workers by machines, and which complements abstract (high-earning) workers. The displaced routine workers find either abstract or manual (low-earning) jobs, increasing the employment share of these occupations (Autor, Levy, and Murnane (2003), Autor, Katz, and Kearney (2006), Michaels, Natraj, and Van Reenen (2014), Autor and Dorn (2013), Goos, Manning, and Salomons (2014)). Autor and Dorn (2013) argue that much of the expansion of manual occupations is driven by the expansion of low-skilled service jobs, our definition of low-end service jobs echo this argument. While the spread of ICT is a convincing explanation for the polarization of labor markets after the mid-1980s, it does not provide an explanation for the patterns observed earlier.[8]

The remainder of the paper is organized as follows: Section 2 lays out our empirical findings, Section 3 our theoretical model, Section 4 the quantitative results, and Section 5 concludes.

[5] A notable exception is Caselli and Coleman (2001).

[6] If ICT is used differentially across sectors, and ICT becomes cheaper, this would show up as an increase in relative labor productivity in the sector which uses ICT most intensively.

[7] Another explanation is the increasing off-shorability of tasks (rather than finished goods), as first emphasized by Grossman and Rossi-Hansberg (2008). It has been argued that it is largely the middle-earning jobs that are off-shorable, but the evidence is mixed (Blinder (2009), Blinder and Krueger (2013), Acemoglu and Autor (2011)).

[8] The same is true for the off-shoring explanation.

2 Polarization in the data

Using US Census data between 1950 and 2000 and the 2007 American Community Survey (ACS), we document the following three facts: 1) polarization in terms of occupations started as early as the 1950/60s, 2) wages and employment have been polarizing in terms of broadly defined industries, 3) a significant part of employment polarization in terms of occupations is driven by employment shifts across industries. In what follows we document each of these facts in detail.

2.1 Polarization in terms of occupations

In the empirical literature, polarization is mostly represented in terms of occupations. We document polarization in terms of three occupational classifications: we start from the finest balanced occupational codes possible, then go to ten broad occupational categories, and finally present polarization in terms of three commonly used categories in the literature.

Figure 1: Wage and employment polarization

Notes: The data is taken from IPUMS US Census data for 1950, 1960, 1970, 1980, 1990, 2000 and the American Community Survey (ACS) for 2007. The sample excludes agricultural occupations/industries and observations with missing wage data; the details are given in the Appendix. Balanced occupation categories (183 of them) were defined by the authors based on Meyer and Osborne (2005), Dorn (2009) and Autor and Dorn (2013). The bottom two panels show the 30-year change in employment shares (calculated as hours supplied rather than persons), and the top two panels show the 30-year change in log hourly real wages (again labor supply weighted). In the left panels occupations are ranked based on their 1950 average wage, whereas in the right panels they are ranked according to their 1980 average wage.

Following the methodology used in Autor, Katz, and Kearney (2006), Acemoglu and Autor (2011), and Autor and Dorn (2013), we plot the smoothed changes in employment shares and log real wages for a balanced panel of occupational categories ranked according to their 1950 and 1980 mean wages. The novelty in these graphs is that we show these patterns going back until 1950, whereas most analyses look at data only from 1980 onwards. The top row of Figure 1 shows that there has been polarization in terms of real wages in all 30-year periods. This polarization is present whether the occupations are ranked according to their 1950 or 1980 mean wages. The polarization of real wages is most pronounced in the first two 30-year intervals, but it is clearly discernible in the following ones as well from the slight U-shape of the smoothed changes. The picture is more mixed in terms of employment polarization (the bottom row of Figure 1): employment polarization is most pronounced in the last 30 years (1980-2007), but it seems to be present even in the earlier decades.[9]

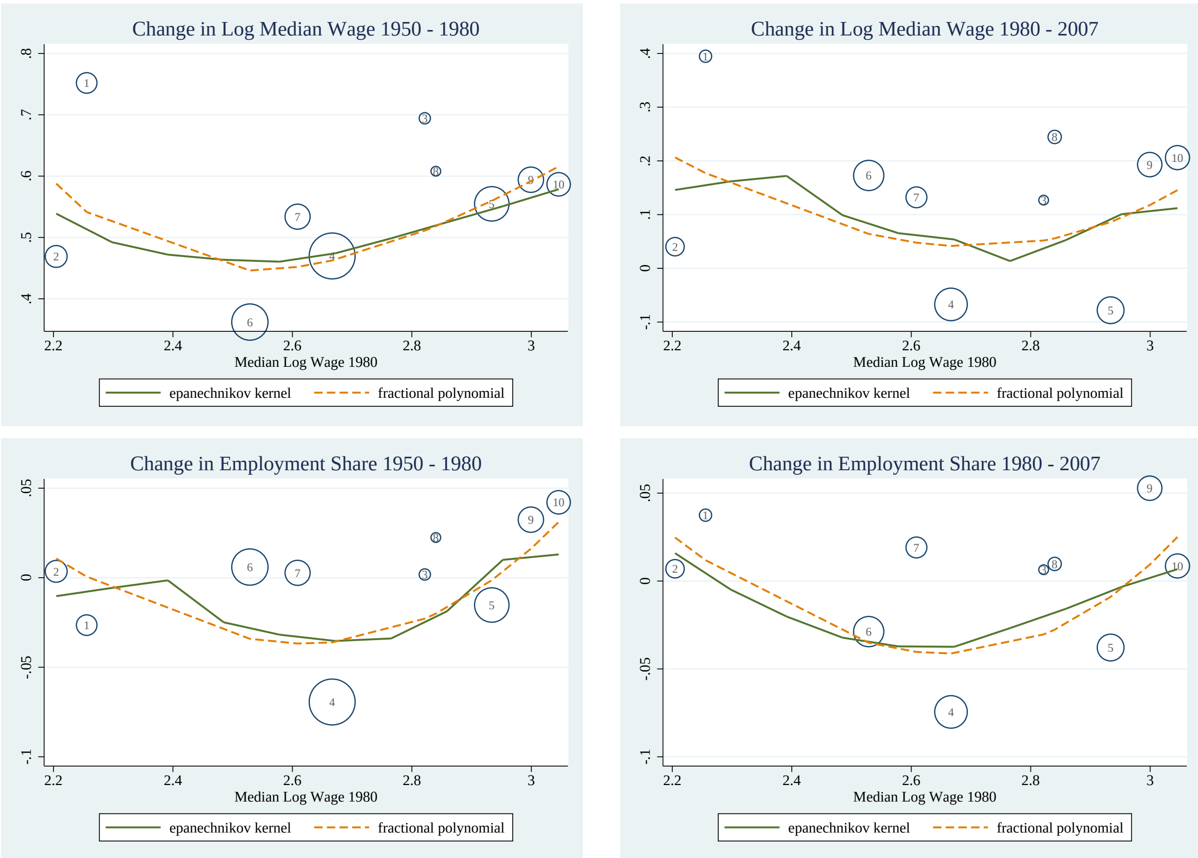

Figure 2: Polarization in broad occupational categories

Notes: These graphs plot for 10 occupational categories the 30-year change in the occupation's median log wage (top row) and in its employment share (bottom row) against their median log wage in 1980. The size of the circles represents the occupations employment share in the initial year. The left panels show the changes between 1950 and 1980, whereas the right show the changes between 1980-2007. The 10 occupational codes are 1 personal care, 2 food and cleaning services, 3 protective services, 4 operators, fabricators and laborers, 5 production, construction trades, extractive and precision production, 6 administrative and support occupations, 7 sales, 8 technicians and related support occupations, 9 professional specialty occupations, 10 managers.

An alternative way to present these trends is to group occupations into broad categories and com-

[9]These patterns do not necessarily hold for a decade-by-decade analysis. In some decades the top gains, whereas in others the bottom gains, but it is never the middle that grows the most in terms of employment shares. See Figure 11 in the Appendix.

pute how their real wages and employment shares have evolved. Similarly to Goos, Mannings, and Salomons (2007 and 2009), Figure 2 shows for 10 occupational groups the change in their median log wages (top row) and in their employment shares (bottom row) against their median log wage in 1980. The panels on the left show the change over 1950-1980 and the panels on the right over 1980-2007. The size of the marker in the scatter plot corresponds to the employment share of the occupation in the initial year. The plots also show two nonlinear fits: the solid line is an epanechnikov kernel and the dashed line is a second order fractional polynomial. This analysis confirms polarization of real wages and employment, starting as early as 1950.

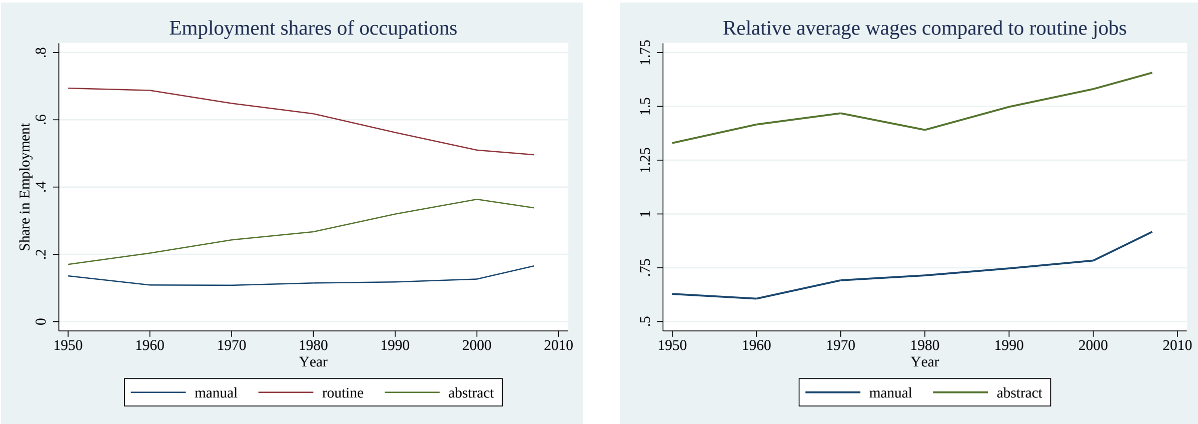

Figure 3: Polarization for broad occupations

Notes: Employment shares (in terms of hours) and relative average wages are calculated from the same data as in Figure 1. For details of the occupation classification see text and the Appendix.

Finally, we document polarization in terms of occupations in an even coarser classification. As in Acemoglu and Autor (2011) we classify occupation groups into the following categories: manual, routine, and abstract.[10] Figure 3 shows the patterns of polarization both in terms of employment shares and wages between 1950 and 2007 for these three broad categories. The left panel shows that the employment share of routine occupations has been falling, of abstract occupations has been increasing since the 1950s, while of manual occupations, following a slight compression until 1960, has been steadily increasing. The right panel shows the path of the relative average manual and abstract wage compared to the routine wage. It is worth to note that, as expected, manual workers on average earn less than routine workers, while abstract workers earn more. However, over time, the advantage of routine jobs over manual jobs has been falling, and the advantage of abstract jobs over routine jobs has been rising. Thus, the middle earning group, the routine workers, lost both in terms of relative average wage and employment share to the benefit of manual and abstract workers. In other words, in terms of these three broad occupations there is clear evidence for polarization.

[10] Details on these classifications are given in the Appendix.

2.2 Polarization in terms of sectors

Next we document the polarization of employment and wages in terms of three broad industries or sectors: low-skilled services, manufacturing, and high-skilled services. As discussed in the introduction, this categorization is driven by both demand-side and production-side considerations. Our manufacturing classification is broader than what is commonly used in the literature, as it includes wholesale trade and retail trade as well.[11] Our rationale for including wholesale and retail trade in manufacturing is that these services are typically provided jointly with a manufacturing good, hence the demand for these are driven by the demand for the good that is being sold. We split the remaining (service) industries into two: low-skilled services are those which could be home produced, while high-skilled services are those which could only be produced on the market.[12]

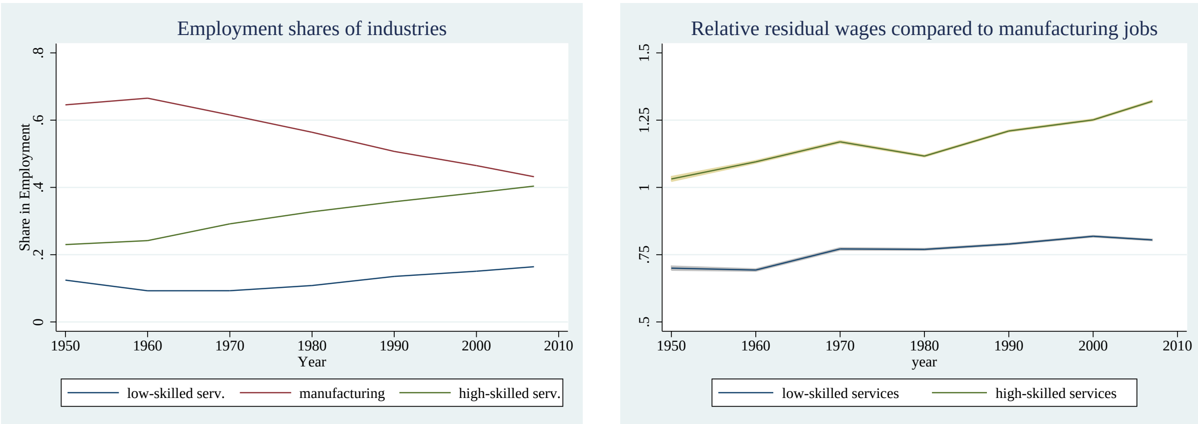

Figure 4: Polarization for broad industries

Notes: The data used is the same as in Figure 1. Each worker is classified into one of three sectors based on the ind1990 code (for details of the industry classification see text and the Appendix.) The left panel shows employment shares, calculated in terms of hours worked. The right panel shows relative wages: the high-skilled service and the low-skilled service premium compared to manufacturing (and their 95% confidence intervals), implied by the regression of log wages on gender, race, a polynomial in potential experience, and sector dummies.

Figure 4 shows the patterns of polarization both in terms of employment shares and wages for the above defined sectors between 1950 and 2007. The left panel shows the path of employment shares: high-skilled services increase continuously, low-skilled services increase and manufacturing decreases from 1960 onwards.[13] In terms of wages, we plot the sector premium in high-skilled and low-skilled services, as well as their 95 percent confidence intervals. These sector premia are the exponents of the coefficients on sector dummies, which come from a regression of log wages where we control also for gender, race, and a polynomial in potential experience.[14] We plot these rather than the relative average

[11]It also includes mining, construction, manufacturing, transport and public utilities, all of which are classified as manufacturing in some of the structural change literature: Buera et al (2012), Kongsamut et al (2001), Ngai et al (2008), Herrendorf et al (2013).

[12]Details on these classifications are given in the Appendix. While we are not using the educational attainment of employees in categorizing industries, the ranking that results from our distinction between substitutes for home production and market services turns out to be such that indeed low-skilled services have a less skilled workforce than high-skilled services. See Table 4 in the Appendix.

[13]Between 1950 and 1960 manufacturing employment was still increasing, and low-skilled services saw a moderate drop.

[14]See Table 5 in the Appendix for details of the regressions. The patterns of relative average wages for industries are very

wages (as in Figure 3 for occupations), because in our quantitative exercise we do not aim to explain sectoral wage differentials that are potentially caused by age, gender or racial composition differences and the differential path of these across sectors.[15] As the graph shows, manufacturing workers are the middle earners, low-skilled service workers have a wage discount, whereas high-skilled service workers have a wage premium compared to them. The premium of high-skilled workers is increasing from 1950 to 2007, while the discount of low-skilled service workers is falling throughout. To summarize, there is clear polarization in terms of these three sectors: the low- and high-earners gained in terms of employment and wages at the expense of the middle-earning, manufacturing workers.

2.3 Polarization across occupations linked to industry employment shifts

To quantify the contribution of sectoral employment shifts to each occupation's employment share path, we conduct a standard shift-share decomposition.[16] The overall change in the employment share of occupation o between year 0 and t, $\Delta E_{ot} = E_{ot} - E_{o0}$, can be expressed as:

$$\Delta E_{ot} = \underbrace{\sum_i \lambda_{oi} \Delta E_{it}}_{\equiv \Delta E_{ot}^B} + \underbrace{\sum_i \Delta \lambda_{oit} E_i}_{\equiv \Delta E_{ot}^W},$$

where ΔE_o^B represents the change in the employment share of occupation o that is attributable to changes in industrial composition, i.e. structural transformation, while ΔE_o^W reflects changes driven by within sector forces. The change driven by shifts between sectors is calculated as the weighted sum of the change in sector i's employment share, ΔE_{it}, where the weights are the average share of occupation o within sector i, $\lambda_{oi} = (\lambda_{oit} + \lambda_{oi0})/2$. The change driven by shifts within sectors is calculated as the weighted sum of the change in occupation o's share within sector i employment, $\Delta\lambda_{oit}$, where the weights are the average employment share of sector i, $E_i = (E_{it} + E_{i0})/2$.[17]

Table 1 shows the decomposition of occupational employment share changes between 1950 and 2007 into a between industry and a within industry component. The first column uses the three occupation categories (manual, routine, abstract) and the three sectors (low-skilled services, manufacturing, high-skilled services) defined earlier. To be sure that our results are not driven by the coarse categorization, similarly to Acemoglu and Autor (2011), we implement this decomposition also in terms of broader categories, 10 occupations and 11 industry groups, shown in column two and four.[18]

similar, see Figure 12 in the Appendix. Also, in terms of occupational wage differentials, the path of occupational premia is similar to the relative average wage paths, see Figure 12 and Table 6 in the Appendix.

[15] One could be concerned that the employment share changes are driven by changes in the age, gender, race composition of the labor force. To control for this, we generate counterfactual industry employment shares, by fixing the industry employment share of each age-gender-race cell at its 1950 level, and allowing the employment shares of the cells to change. This exercise confirms that the employment share changes are not driven by the compositional changes of the labor force. See Figure 13 in the Appendix.

[16] An alternative way is to calculate how much occupational employment shares would have changed, if industry employment shares would have remained at their 1950 level. See Figure 14 in the Appendix.

[17] $\lambda_{oit} = L_{oit}/L_{it}$ is the share of occupation o, industry i employment within industry i employment at time t. $E_{it} = L_{it}/L_t$ is the share of industry i employment within total employment at time t.

[18] The 10 occupations are the same as in Figure 2, while the 11 industries are: 1 business and service repairs, 2 personal services,

Table 1: Decomposition of changes in occupational employment shares over 1950-2007

	Employment shares	
	3 x 3	10 x 11
Manual		
Total Δ	2.98	3.12
Between Δ	4.39	3.98
Within Δ	-1.41	-0.85
Routine		
Total Δ	-19.79	-25.80
Between Δ	-10.46	-12.38
Within Δ	-9.33	-13.42
Abstract		
Total Δ	16.81	19.79
Between Δ	6.07	8.94
Within Δ	10.74	10.84
Average		
Total Δ	-21.16	-20.04
Between Δ	-12.05	-8.83
Within Δ	-9.10	-11.21

Notes: Same data as in Figure 1. For each occupational category, the first row presents the total change, the second the between industry component, and the third the within industry component between 1950 and 2007. The first column uses 3 occupations and 3 sectors, the second 10 occupations and 11 industries.

The polarization in terms of these three broad occupational categories documented in Figure 3 is confirmed by Table 1. The employment share of manual occupations increased by 3 percentage points, of abstract occupations it increased by 17 percentage points, and in routine occupations it contracted by 20 percentage points.

In terms of average occupational employment share changes between 44 per cent and 56 per cent of the changes are driven by between industry shifts, depending on whether we consider the finer or the coarser categories. For manual occupations, it seems that the within-industry shifts would have lead to a decline in its employment share, and hence all of the increase is driven by between industry factors. For routine occupations between 48 and 53 percent of the fall in employment share is due to between industry shifts. For abstract occupations the role of between industry shifts is slightly less important, it explains between 36 and 45 percent of the total change. This decomposition indicates that a significant part of the occupational employment share changes are driven by shifts in the industrial composition of the economy between 1950 and 2007. The role of between industry shifts are similar in magnitude to those found by Acemoglu and Autor (2011) and Goos, Manning, and Salomons (2014). Acemoglu and Autor (2011) use US Census data between 1960 and 2000, and the ACS 2008 data to decompose occupational employment share changes. Their focus is the declining importance of between industry shifts from 1960-1980 to 1980-2007. In our decomposition, we do not find strong evidence for the diminished

nursing and personal care facilities, child day care services, family child care homes, residential care facilties, taxi, retail bakeries, eating and drinking places, entertainment, 3 extractive industries, 4 construction, 5 manufacturing, 6 transport and public utilities, 7 wholesale trade, 8 retail trade, 9 finance, insurance, and real estate, 10 professional and related services, advertising, computer and data processing services, communications, theaters and motion pictures, 11 public administration.

role of between industry shifts, this might be due to the fact that we do not look at men and women separately, or that we look at hours worked rather than raw employment numbers. Since in our general equilibrium model, labor demand is an important factor, our focus is the change in overall demand, rather than demand for male and female labor separately. Goos, Manning, and Salomons (2014) use data for 16 European countries between 1993 and 2010, and attribute a roughly equal role to between and within industry shifts in all occupations. Goos et al (2014) argue that part of the between industry shifts can be driven by routinization, which is a within industry phenomenon. Since routinization has a bigger impact on industries where routine labor is used more intensively, employment might shift away from these industries. We stay agnostic about the origin of the faster technological progress in manufacturing, and indeed it is possible that since the 1980s or 1990s part of the faster growth is due to the more intensive use of routine labor combined with routinization. However, routinization, driven by ICT, is not likely to be driving the faster productivity growth observed between the 1950s and 1980s.

Our reading of this decomposition is that in order to understand the occupational employment share changes it is important to consider the forces that drive the structural shift of employment away from manufacturing and towards both types of services.

To summarize: we documented three new facts about the occupational and sectoral employment shares and relative wages. We documented that polarization defined over occupational categories (finer and broader) both in terms of employment and wages has been present in the US since the 1950s. Second we showed that the same patterns are discernible in terms of three broad sectors: low-skilled services, manufacturing and high-skilled services. Finally, we showed that a significant amount of the employment share changes in occupations is driven by the employment shifts across industries.

In the rest of the paper we present a model to jointly explain the sectoral shifts in employment and the changes in the sectoral wage premia. We then calibrate the model and quantitatively assess how much of the sectoral patterns can it explain between 1950 and 2007.

3 Model

In order to illustrate the mechanism that is driving the polarization of wages and employment, we present a parsimonious static model, and analyze its behavior as productivity levels increase across sectors. The key novel feature of our model is that we assume that each sector values different skills in its production process. Relaxing the assumption of the homogeneity of labor allows us to derive predictions, not only about the employment and expenditure shares, but also about the relative average wages across sectors over time.

We assume that the economy is populated by heterogenous agents, who all make individually optimal decisions about their sector of work. Every individual chooses their sector of work to maximize earnings, in a Roy-model type setup. We assume that in low-skilled services everyone is equally productive, as everyone uses the one unit of raw labor that they have. On the other hand, we assume that individuals are ex ante heterogeneous in their efficiency units of labor in manufacturing and in high-skilled services, and thus endogenously sort into the sector where the return to their labor is the highest.

Furthermore these individuals are organized into a stand-in household, which collects all earnings, and provides the same consumption basket for each of its members. The household maximizes its utility subject to the household level budget constraint. Households derive utility from consuming high- and low-skilled services and manufacturing goods.

The economy is in a decentralized equilibrium at all times: individuals make sectoral choices to maximize their earnings, the stand-in household collects all earnings and maximizes its utility by optimally allocating this income between low-skilled services, manufacturing goods and high-skilled services. Production is perfectly competitive, wages and prices are such that all markets clear. We analyze the qualitative and quantitative role of technological progress in explaining the observed wage and employment dynamics since the 1950s.

3.1 Sectors and production

There are three sectors in the model: high-skilled services (S), manufacturing (M), and low-skilled services (L). All goods and services are produced in perfect competition, and each sector uses only labor as an input into production.

The technology to produce high-skilled services is:

$$Y_s = A_s N_s, \tag{1}$$

where A_s is productivity and N_s is the total amount of efficiency units of labor hired in sector S for production. Sector S firms are price takers, therefore the equilibrium wage per efficiency unit of labor has to satisfy:

$$w_s = \frac{\partial p_s Y_s}{\partial N_s} = p_s A_s. \tag{2}$$

The technology to produce manufacturing goods is:

$$Y_m = A_m N_m, \tag{3}$$

where A_m is productivity, N_m is the total amount of efficiency units of labor hired in sector M. Since sector M firms are also price takers, the equilibrium wage per efficiency unit of labor in sector M has to satisfy:

$$w_m = \frac{\partial p_m Y_m}{\partial N_m} = p_m A_m. \tag{4}$$

Note that the wage of a worker with a efficiency units of sector i labor working in sector $i \in \{M, S\}$ is $w_i a$.

We assume that each worker is equally talented in providing low-skilled services, i.e. efficiency units of labor do not matter here, only the raw amount of hours that a worker can provide is important. The total amount of low-skilled services provided is given by:

$$Y_l = A_l L_l, \tag{5}$$

where A_l is low-skilled service productivity, and L_l is the raw units of labor working in the low-skilled service sector. Since sector L firms are also price takes, the equilibrium wage per unit of raw labor in sector L has to satisfy:

$$w_l = \frac{\partial p_l Y_l}{\partial L_l} = p_l A_l. \tag{6}$$

Note that since everyone has the same amount of raw labor, this implies that everyone working in the low-skilled service sector has the same earnings.

3.2 Labor supply and demand for goods

The stand-in household consists of a measure one continuum of different types of members. Each member chooses which one of the three market sectors to supply his one unit of raw labor in. The household collects the earnings of all its members and decides how much low-skilled services, manufacturing goods and high-skilled services to buy on the market.

3.2.1 Sector of work

We assume that every member of the households works full time in one of the three market sectors. Since every member can work in any of the three sectors, and each member's utility is increasing in his own earnings (as well as in all other members' earnings), it is optimal for each worker to choose the sector which provides him with the highest earnings.

Individuals are heterogeneous in their innate ability, $\mathbf{a} \in \mathbf{R}^2_+$, which is drawn from a time invariant distribution $f(\mathbf{a})$. For simplicity we assume that $a_m \equiv \mathbf{a}(1)$ denotes the individual's efficiency units of labor in manufacturing, while $a_s \equiv \mathbf{a}(2)$ denotes his efficiency units in high-skilled services.[19] We further assume that each individual is equally productive when working in low-skilled services, as that only requires an individual's raw labor. Therefore the earnings of an individual with innate ability $\mathbf{a} = (a_m, a_s)$ in sector L is w_l, in sector M it is $a_m w_m$, while if working in sector S it is $a_s w_s$.

Given wage rates w_l, w_m, w_s – per raw unit of labor in L and per efficiency unit in sector M and S – the optimal decision of any agent can be characterized as follows.

Result 1. *Given wage rates w_l, w_m and w_s, the optimal sector choice of individuals can be characterized by two*

[19] The results of the model are qualitatively unchanged if we assume that an individual has $\alpha_m \mathbf{a}(1) + (1 - \alpha_m)\mathbf{a}(2)$ efficiency units of labor in manufacturing, while he has $\alpha_s \mathbf{a}(1) + (1 - \alpha_s)\mathbf{a}(2)$ in high-skilled services, as long as $\alpha_m \neq \alpha_s$.

cutoff values:

$$\widehat{a}_m \equiv \frac{w_l}{w_m} \tag{7}$$

$$\widehat{a}_s \equiv \frac{w_l}{w_s}. \tag{8}$$

It is optimal for an individual with innate ability (a_m, a_s) to work in sector L if and only if

$$a_m \leq \widehat{a}_m \quad \textit{and} \quad a_s \leq \widehat{a}_s. \tag{9}$$

It is optimal for the individual to work in sector M if and only if

$$a_m \geq \widehat{a}_m \quad \textit{and} \quad a_s \leq \frac{\widehat{a}_s}{\widehat{a}_m} a_m. \tag{10}$$

Finally it is optimal to work in sector S if and only if

$$a_s \geq \widehat{a}_s \quad \textit{and} \quad a_m \leq \frac{\widehat{a}_m}{\widehat{a}_s} a_s \tag{11}$$

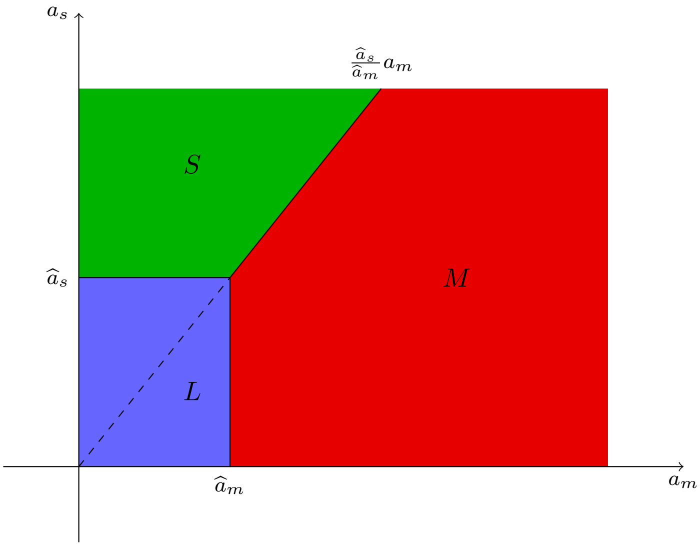

Figure 5: Optimal sector of work

The figure depicts the optimal sector of work choices as a function of $\widehat{a}_m = w_l/w_m$ and $\widehat{a}_s = w_l/w_s$. The blue area shows the efficiency unit pairs (a_m, a_s) where L is the optimal sector, the red area shows where M is optimal, and the green area shows where S is optimal.

Figure 5 shows this endogenous sorting behavior. Individuals who have low efficiency units or low productivity in both manufacturing and high-skilled services sort into low-skilled services (the blue area). Individuals with high enough manufacturing efficiency and relative to this a low high-skilled service efficiency sort into manufacturing jobs (the red area). While those individuals who have a high enough high-skilled service efficiency and relative to this a low manufacturing efficiency choose to work in high-skilled services (the green area).

Given the optimal sector of work choices of individuals the effective labor supplies in the three markets are:

$$L_l(\widehat{a}_m, \widehat{a}_s) = \int_0^{\widehat{a}_m} \int_0^{\widehat{a}_s} f(a_m, a_s) da_s da_m, \tag{12}$$

$$N_m(\widehat{a}_m, \widehat{a}_s) = \int_{\widehat{a}_m}^{\infty} \int_0^{\frac{\widehat{a}_s}{\widehat{a}_m} a_m} a_m f(a_m, a_s) da_s da_m, \tag{13}$$

$$N_s(\widehat{a}_m, \widehat{a}_s) = \int_{\widehat{a}_s}^{\infty} \int_0^{\frac{\widehat{a}_m}{\widehat{a}_s} a_s} a_s f(a_m, a_s) da_m da_s. \tag{14}$$

Note that these are the effective labor supplies in sector M and S, the raw labor supplies in these sectors are:

$$L_m(\widehat{a}_m, \widehat{a}_s) = \int_{\widehat{a}_m}^{\infty} \int_0^{\frac{\widehat{a}_s}{\widehat{a}_m} a_m} f(a_m, a_s) da_s da_m, \tag{15}$$

$$L_s(\widehat{a}_m, \widehat{a}_s) = \int_{\widehat{a}_s}^{\infty} \int_0^{\frac{\widehat{a}_m}{\widehat{a}_s} a_s} f(a_m, a_s) da_m da_s. \tag{16}$$

It is worth to note that ceteris paribus a higher $\widehat{a}_m$, i.e. a higher L unit wage relative to M unit wage, leads more people to choose sector L and sector S over M, leading to an increase L_l and N_s and a reduction in N_m. A higher $\widehat{a}_s$, i.e. a higher unit wage in L relative to S induces more people to select sector L and M over sector S, increasing L_l and N_m, and reducing N_s. An increase in $\widehat{a}_m/\widehat{a}_s$ incentivizes more people to choose sector S over M, while its impact on the L vs M and L vs S choice depends on the absolute change in the cutoffs.

3.2.2 Demand for consumption goods and services

Household members derive utility from low-skilled services, manufacturing goods and high-skilled services. The household allocates total income earned by household members to maximize the following utility:

$$\max_{C_l, C_m, C_s} \ln \left(\theta_l C_l^{\frac{\varepsilon-1}{\varepsilon}} + \theta_m C_m^{\frac{\varepsilon-1}{\varepsilon}} + \theta_s C_s^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}$$

$$\text{s.t.} \quad p_l C_l + p_m C_m + p_s C_s \leq w_l L_l + w_m N_m + w_s N_s \equiv m$$

where m is the total earnings of the household, p_l, p_m, and p_s are the prices of the low-skilled services, the manufacturing goods, and the high-skilled services.

The household's optimal consumption bundle has to satisfy:

$$\frac{C_l}{C_m} = \left(\frac{p_l}{p_m} \frac{\theta_m}{\theta_l} \right)^{-\varepsilon}, \tag{17}$$

$$\frac{C_s}{C_m} = \left(\frac{p_s}{p_m} \frac{\theta_m}{\theta_s} \right)^{-\varepsilon}. \tag{18}$$

3.3 Competitive equilibrium and structural change

A competitive equilibrium is given by cutoff sector of work abilities $\{\widehat{a}_m, \widehat{a}_s\}$, wage rates$\{w_l, w_m, w_s\}$, prices $\{p_l, p_m, p_s\}$, and consumption demands $\{C_l, C_m, C_s\}$, given productivities $\{A_l, A_m, A_s\}$, where individuals, households and firms make optimal decisions and all markets clear.

Using goods market clearing in all sectors ($Y_i = C_i$ for $i = L, M, S$), market clearing wage rates,(2), (4) and (6), in the household's optimality conditions, (17) and (18), we obtain the following:

$$\frac{A_l}{A_m}\frac{L_l(\widehat{a}_m, \widehat{a}_s)}{N_m(\widehat{a}_m, \widehat{a}_s)} = \left(\frac{w_l}{w_m}\frac{A_m}{A_l}\frac{\theta_m}{\theta_l}\right)^{-\varepsilon},$$
$$\frac{A_s}{A_m}\frac{N_s(\widehat{a}_m, \widehat{a}_s)}{N_m(\widehat{a}_m, \widehat{a}_s)} = \left(\frac{w_s}{w_m}\frac{A_m}{A_s}\frac{\theta_m}{\theta_s}\right)^{-\varepsilon}.$$

The above shows that an increase in relative manufacturing productivity compared to low-skilled service productivity (A_m/A_l) has two opposing direct effects: an increase in the relative price of low-skilled services (p_l/p_m), and a fall in relative low-skilled service output, (Y_l/Y_m). While the increase in relative price leads to a fall in relative demand for low-skilled services, there is a contemporaneous fall in relative supply of low-skilled services. If low-skilled services and manufacturing goods are complements, $\varepsilon < 1$, the relative price has a smaller impact, and the relative effective employment in low-skilled services has to increase for equilibrium, which requires a rise in relative wages, w_l/w_m. The effects of an increase in A_m/A_s similarly leads to an increase in the relative effective high-skilled service employment compared to manufacturing, which requires a rise in w_s/w_m.

To formally analyze the comparative static properties of the equilibrium, we use the optimal sector of work cutoffs, (29) and (28), in the above two equations. Re-arranging we obtain the following expressions:

$$\frac{L_l(\widehat{a}_m, \widehat{a}_s)}{N_m(\widehat{a}_m, \widehat{a}_s)}\widehat{a}_m^{\varepsilon} = \left(\frac{A_m}{A_l}\right)^{1-\varepsilon}\left(\frac{\theta_m}{\theta_l}\right)^{-\varepsilon}, \tag{19}$$
$$\frac{N_s(\widehat{a}_m, \widehat{a}_s)}{N_m(\widehat{a}_m, \widehat{a}_s)}\left(\frac{\widehat{a}_m}{\widehat{a}_s}\right)^{\varepsilon} = \left(\frac{A_m}{A_s}\right)^{1-\varepsilon}\left(\frac{\theta_m}{\theta_s}\right)^{-\varepsilon}. \tag{20}$$

These two equations implicitly define the equilibrium sector of work cutoffs, $\widehat{a}_m$ and $\widehat{a}_s$, and in turn these cutoffs fully characterize the equilibrium of the economy.

Proposition 1. *When manufacturing goods and the two types of services are complements ($\varepsilon < 1$), then faster productivity growth in manufacturing than in both types of services ($dA_m/A_m > dA_s/A_s = dA_l/A_l$), leads to a change in the optimal sorting of individuals across sectors. In particular $\widehat{a}_m$ unambiguously increases, $\widehat{a}_s$ can rise or fall, but $d\widehat{a}_m/\widehat{a}_m > d\widehat{a}_s/\widehat{a}_s$. This results in an unambiguous increase in employment in L, in effective employment in S, and a reduction in effective and raw employment in M.*

Proof. Total differentiation of (19) and (20). See Appendix for details. □

Proposition 1 confirms the results of Ngai and Pissarides (2007) in terms of efficiency labor, rather than hours worked: when sectoral outputs are complements in consumption, efficiency labor needs to reallocate to the sectors which become relatively less productive. Since in low-skilled services efficiency and raw labor units are the same, the observed employment share of low-skilled services should rise. To understand how raw employment changes in manufacturing and high-skilled services Figure 6 is useful. This figure shows the possible configurations of responses in optimal sector of work cutoffs. Proposition 1 states that the boundary between L and M shifts out ($\widehat{a}_m$ increases), while the boundary between S and M becomes flatter ($\widehat{a}_s/\widehat{a}_m$ decreases). It is ambiguous in general whether the boundary between L and S shifts up or down ($\widehat{a}_s$ increases or decreases). The left panel in Figure 6 shows the case when $\widehat{a}_s$ decreases, while the right panel shows the case when it increases.

Effective employment in M is the aggregate amount of a_m in the area bounded by $\widehat{a}_m$ and $a_m\widehat{a}_s/\widehat{a}_m$, while raw employment is just the mass of individuals in the same area. Both the outward shift of $\widehat{a}_m$ and the flattening of $\widehat{a}_s/\widehat{a}_m$ have a negative impact on the raw and effective employment in M.

Similarly effective employment in S is the aggregate amount of a_s in the area bounded by $\widehat{a}_s$ and $a_m\widehat{a}_s/\widehat{a}_m$, while raw employment is the mass in the same area. If $\widehat{a}_s$ falls, as shown in the left panel, then effective and raw employment in S increases. If $\widehat{a}_s$ increases, shown in the right panel, raw employment is S might fall, while effective employment in S has to increase. The fact that effective employment has to increase implies that the amount of a_s lost due to the increase in $\widehat{a}_s$ is smaller than the amount of a_s gained due to the decrease in $\widehat{a}_s/\widehat{a}_m$. However, since the average a_s in the area lost is strictly lower than in the area gained, we cannot determine whether raw employment increases or falls in S.

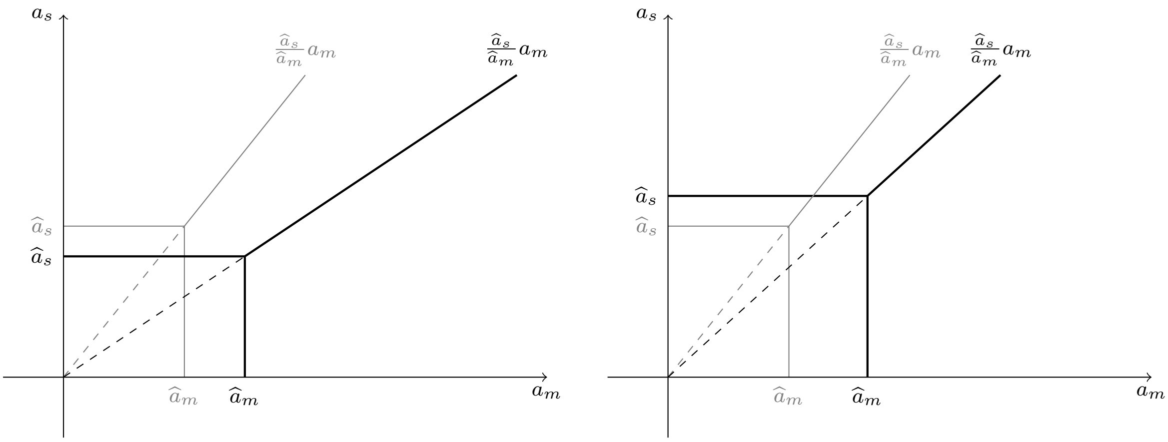

Figure 6: Change in the optimal sorting

The figure shows the possible shifts in $\widehat{a}_m$ and $\widehat{a}_s$ in response to an increase in relative manufacturing productivity when the sectoral outputs are complements in consumption. While $\widehat{a}_m$ unambiguously increases, and $\widehat{a}_s/\widehat{a}_m$ unambiguously decreases, the direction of change in $\widehat{a}_s$ is ambiguous.

The model's predictions in terms of relative average wages are not clear cut. To see this, consider the average low-skilled service wage relative to the average manufacturing wage:

$$\frac{\overline{w}_l}{\overline{w}_m} = \frac{w_l}{w_m \frac{N_m}{L_m}} = \frac{\widehat{a}_m}{\overline{a}_m},$$

where $\overline{a}_m \equiv N_m/L_m$ is the average efficiency of manufacturing workers in manufacturing. The percentage change in the average wage of low-skilled service workers relative to manufacturing workers is just:

$$\frac{d\frac{\overline{w}_l}{\overline{w}_m}}{\frac{\overline{w}_l}{\overline{w}_m}} = \frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\overline{a}_m}{\overline{a}_m},$$

i.e. the difference between the percentage change in the cutoff ability between L and M and the percentage change in the average ability in sector M. From Proposition 1 we know that $\frac{d\widehat{a}_m}{\widehat{a}_m}$ is positive, the question is whether the average ability in M can increase more in percentage terms than this. As long as the raw employment in M is not very small, and the distribution of abilities is relatively smooth, the relative average wage in low-skilled services will increase.

Similarly the high-skilled service wage relative to the average manufacturing wage can be expressed as:

$$\frac{\overline{w}_s}{\overline{w}_m} = \frac{w_s\frac{N_s}{L_s}}{w_m\frac{N_m}{L_m}} = \frac{\widehat{a}_m}{\widehat{a}_s}\frac{\overline{a}_s}{\overline{a}_m}.$$

Implying that the percentage change in the relative wage is:

$$\frac{d\frac{\overline{w}_s}{\overline{w}_m}}{\frac{\overline{w}_s}{\overline{w}_m}} = \frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s} + \frac{d\overline{a}_s}{\overline{a}_s} - \frac{d\overline{a}_m}{\overline{a}_m}.$$

Again this change cannot be signed independent of the underlying distribution of abilities and of the initial level of employment in the different sectors.

We can show that relative sectoral value added shares increase in the sectors with lower productivity growth if the sectoral outputs are complements in consumption.

Proposition 2. *When manufacturing goods and the two types of services are complements ($\varepsilon < 1$), then faster productivity growth in manufacturing than in both types of services ($dA_m/A_m > dA_s/A_s = dA_l/A_l$), increases the relative value added in both high- and low-skilled services compared to manufacturing:*

$$d\frac{p_sY_s}{p_mY_m} > 0 \quad \textit{and} \quad d\frac{p_lY_l}{p_mY_m} > 0.$$

These results can be seen from the following. In this model, sectoral value added is equal to the sectoral wage bill: $p_iY_i = p_iA_iN_i = w_iN_i$. Proposition 1 tells us that $w_l/w_m = \widehat{a}_m$ increases, that L_l increases, and N_m falls. Both relative unit wages and effective employment changes increase the value added output of sector L relative to sector M. Similarly, $w_s/w_m = \widehat{a}_m/\widehat{a}_s$ also increases according to Proposition 1, while effective employment in S increases and in M it falls.

The sectoral value added can be further expressed as $p_iY_i = w_iN_i = \overline{w}_iL_i$, since the sectoral wage bill can be expressed as either sectoral unit wage times sectoral effective employment, or as sectoral average wage times sectoral raw employment. Using this latter expression we can show that

$$\frac{p_iY_i}{p_jY_j} = \frac{\overline{w}_i}{\overline{w}_j}\frac{L_i}{L_j}.$$

According to our model relative sectoral value added has to equal the product of relative sectoral average wages and sectoral relative employment. This result holds even if we include capital in the model. In order to break this relationship one needs to assume either imperfect capital mobility across sectors, or different sectoral capital intensities. Since this relationship does not hold in the data, in our calibration we will target relative average wages and sectoral employment shares, as it is the evolution of these two measures that is the focus of our paper.

4 Quantitative results

In this section we quantitatively assess the contribution of structural transformation to the polarization of employment and wages across industries. To do this we consider the evolution of the competitive equilibrium in terms of employment shares and relative average sectoral wages as productivity increases in manufacturing and in both low- and high-skilled services. We calibrate our parameters to match four key moments in 1950, and then feed in the exogenous process for labor productivity to generate predictions for the evolution of employment and wages. We first describe the data targets and the calibration strategy, and then discuss the quantitative importance of our mechanism.

4.1 Calibration

The four key moments are the relative average industry wages, and the industry employment shares. Data for the average industry wages and the industry employment shares come from the 1950 US Census data. Employment shares are calculated as share of hours worked, and relative average wages are the sector premia, both as in section 2.2.

All parameters are time-invariant, and the only exogenous change over time is labor productivity growth. The following parameters need to be calibrated: the parameters of the utility function $\theta_l, \theta_m, \theta_s, \varepsilon$, and the distribution of abilities, $f(a_m, a_s)$.

We assume that the abilities are distributed uniformly. We normalize the mean of a_m and a_s to be unity, as these means cannot be separately identified from other parameters of the model. Given these assumptions on the distribution two parameters are left to be calibrated, the minimum (and thus the maximum, given that the mean is one) of both abilities, denoted by $\underline{a}_m$ and $\underline{a}_s$. We calibrate these two parameters to guarantee that we can match both the raw employment shares and the relative average wages in 1950. The procedure that we implement is the following. For any $\underline{a}_m$ and $\underline{a}_s$, we can calculate the cutoffs $\widehat{a}_m$ and $\widehat{a}_s$, which lead to the raw labor shares seen in the data, i.e. the $L_l(\widehat{a}_m, \widehat{a}_s)$, $L_m(\widehat{a}_m, \widehat{a}_s)$, $L_s(\widehat{a}_m, \widehat{a}_s)$ implied by the cutoffs should match the data. Given $\underline{a}_m, \underline{a}_s$ and $\widehat{a}_m, \widehat{a}_s$, the implied relative average wages can be calculated as follows:

$$\frac{\overline{w}_l}{\overline{w}_m} = \frac{w_l}{w_m \frac{N_m(\widehat{a}_m,\widehat{a}_s)}{L_m(\widehat{a}_m,\widehat{a}_s)}} = \widehat{a}_m \frac{L_m(\widehat{a}_m,\widehat{a}_s)}{N_m(\widehat{a}_m,\widehat{a}_s)},$$

$$\frac{\overline{w}_s}{\overline{w}_m} = \frac{w_s \frac{N_s}{L_s}}{w_m \frac{N_m}{L_m}} = \frac{\widehat{a}_m}{\widehat{a}_s} \frac{L_m(\widehat{a}_m,\widehat{a}_s) N_m(\widehat{a}_m,\widehat{a}_s)}{L_s(\widehat{a}_m,\widehat{a}_s) N_s(\widehat{a}_m,\widehat{a}_s)}.$$

We then calibrate $\underline{a}_m$ and $\underline{a}_s$ to equalize the model implied relative average wages to the ones observed in the data. This procedure guarantees that if we match either the relative average wages or the raw employment shares in 1950, then the other targets are matched automatically; it also allows us to target the ability cutoffs needed to match both the employment shares and the relative average wages.

The parameters of the utility function are left to calibrate. Previous literature has found a very low elasticity of substitution between goods and services when output is measured in value added terms. Ngai and Pissarides (2008) find that plausible estimates are in the range $(0, 0.3)$, while Herrendorf, Rogerson, and Valentinyi (2013) find a value of $\varepsilon = 0.002$, which we use in our baseline calibration. We are thus left to calibrate 3 parameters of the utility function: $\theta_l, \theta_m, \theta_s$. As can be seen from (19) and (20), only the ratio of the θs to the power of $-\varepsilon$ matter. Thus we calibrate $\tau_l \equiv \left(\frac{\theta_m}{\theta_l}\right)^{-\varepsilon}$ and $\tau_s \equiv \left(\frac{\theta_m}{\theta_s}\right)^{-\varepsilon}$ to match the relative average wages in 1950. Without loss of generality we set the initial productivity levels to 1 in each sector.[20] The summary of the calibrated parameters are in the Table 2.

Table 2: Calibrated Parameters

	Description	Value
$[\underline{a}_m, \overline{a}_m]$	range of manufacturing efficiency	$[0.6292, 1.3708]$
$[\underline{a}_s, \overline{a}_s]$	range of high-skilled service efficiency	$[0, 2]$
ε	CES b/w L, M and S in consumption	0.002
τ_l	relative weight on M	0.1752
τ_s	relative weight on S	0.5424

The ability distribution is calibrated to *allow* to jointly match the observed sectoral employment shares and sectoral relative average wages in 1950. Then, conditional on the ability distribution and the elasticity of substitution, ε, we calibrate τ_l and τ_s to *actually* match the employment shares and relative average wages. In the robustness check, we recalibrate the τ's for different values of ε.

Similarly to Ngai and Petrongolo (2014) we calculate labor productivity growth by dividing sectoral value added output data from Herrendorf, Rogerson, and Valentinyi (2013) with sectoral employment data from the Bureau of Economic Analysis (BEA).[21] We rely on Herrendorf, Rogerson, and Valentinyi (2013) for the value added data rather than taking these data directly from the BEA, since as Herrendorf, Rogerson, and Valentinyi point out, the BEA data come from the production side and thus contain both consumption and investment. The authors argue that rather than assuming that all investment is done in manufacturing, the consumption component of each sector's value added needs to be properly calculated. Since the industry level data on value added output and employment is not detailed enough,

[20]We could subsume $\left(\frac{A_m(0)}{A_l(0)}\right)^{1-\varepsilon}$ in τ_l and $\left(\frac{A_m(0)}{A_s(0)}\right)^{1-\varepsilon}$ in τ_s for the calibration. The results would be identical.

[21]We cannot take the growth rates calculated by Ngai and Petrongolo (2014), as our manufacturing classification is wider.

we cannot break down the labor productivity growth of services into low- and high-skilled services. Therefore we assume that productivity growth in low- and high-skilled services is the same. [22]

Table 3: Annual average labor productivity growth

	Based on raw labor		Adjusted by average efficiency	
	Manufacturing	**Services**	**Manufacturing**	**Services**
1950-1960	1.0199	1.0105	1.0213	1.0114
1960-1970	1.0213	1.0103	1.0204	1.0115
1970-1980	1.0115	1.0086	1.0103	1.0089
1980-1990	1.0260	1.0002	1.0242	1.0003
1990-2000	1.0345	1.0068	1.0333	1.0073
2000-2007	1.0229	1.0122	1.0215	1.0128
1950-2007	**1.0226**	**1.0079**	**1.0218**	**1.0085**

Labor productivity growth rates in the first two columns are calculated by dividing sectoral value added data taken from Herrendorf, Rogerson, and Valentinyi (2013) aggregated to our classification of sectors, by sectoral raw employment growth data taken from the BEA. In the second two columns we divide by effective employment growth, which we calculate based on our calibrated distribution of abilities.

Table 3 contains the average annual labor productivity growth in manufacturing and in low- and high-skilled services jointly for each decade between 1950 and 2007, as well as for the entire period. According to our calculations the growth of labor productivity in manufacturing was higher than in the combined services category in each of the decades considered. It is worth to note that both productivity growth and the relative productivity growth between manufacturing and services varied significantly decade by decade. For this reason, for the quantitative evaluation of the model, we first feed in the average growth rates for the entire period, and then the decennial growth rates. The first two columns contain these growth rates by taking raw employment growth.

It is well known that if individuals self-select based on their innate ability and one cannot observe these abilities, then the measurement of changes in average wages or in labor productivity will be biased.[23] In our model, expanding sectors soak up relatively less efficient workers, while contracting sectors shed relatively less efficient workers. This implies that the average ability in manufacturing will increase, while the average ability in high-skilled services will fall, which – if left uncorrected – leads to an overestimation of productivity growth in manufacturing relative to services.[24] To understand the potential magnitude of this bias we calculate the change in average sectoral abilities implied by the change in the cutoff abilities required to match sectoral raw employment shares between 1950 and 2007, given the distribution of underlying abilities. In Table 3 the last two columns show the labor productivity growth corrected for the calculated change in average efficiencies. As it is clear from this table, under our calibration of the ability distribution, the bias in the relative productivity differentials is very small.[25]

[22]See Appendix for details.

[23]Carneiro and Lee (2011) estimated the bias in the measurement of the skill premium, while Young (2014) pointed this out in the context of measuring productivity growth differentials across sectors.

[24]This is true only when manufacturing employment is shrinking, and services employment is expanding.

[25]Young (2014) finds that the implied bias might potentially be so large as to overturn the conventional wisdom of faster productivity growth in manufacturing. However, with our calibrated distribution of abilities the bias is relatively small.

4.2 Wage and employment dynamics

To understand the strength of the mechanisms that we highlight, we simulate the competitive equilibrium of the economy at different productivity levels. We fix the preference and ability distribution parameters at the values calibrated using data only from 1950, and feed in various labor productivity growth measures as calculated in Table 3. Our ultimate interest is the endogenous path of employment and relative average wages.

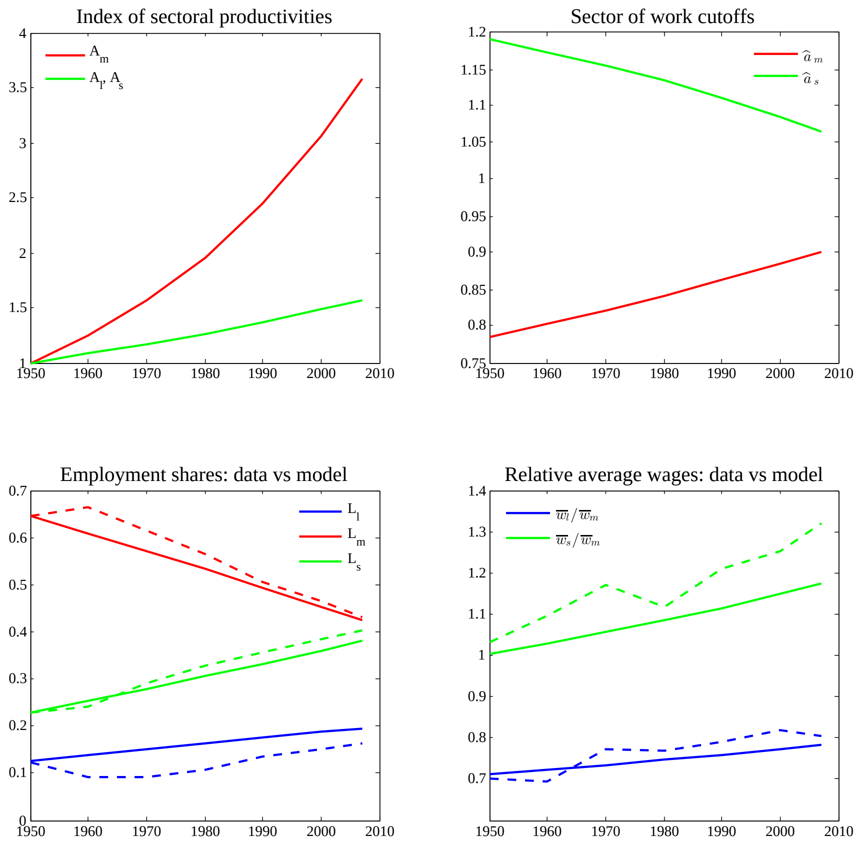

Figure 7: Transition of the benchmark model

The top left panel shows the exogenous change in labor productivity. The top right panel shows the endogenous response of the optimal sector of work cutoffs. The bottom left panel shows the predicted employment shares (solid line) and their actual path (dashed line), while the bottom right panel shows the model predicted (solid line) and actual path (dashed line) of the relative average sectoral wages.

Figure 7 plots the dynamics for our baseline calibration and baseline productivity growth rates. These rates are the annual average raw labor productivity growth for the period 1950-2007: 2.26% annual growth in manufacturing, and 0.79% growth in combined services (bottom left numbers in Table 3). The top left panel shows the path of M and of both L and S sector productivity, we assume that these latter two grow at the same rate. Since productivity growth is highest in the manufacturing sector, but manufactured goods and both types of services are complements in consumption, the increased demand for the output of all sectors in equilibrium is met through a reallocation of labor towards low- and high-skilled services, as we showed in Proposition 1. The increased demand for low- and high-skilled service sector labor puts an upward pressure on their unit wages relative to the unit wage in manufacturing, thus changing the optimal sector of work cutoffs for individuals. The top right panel shows the endogenous response of $\widehat{a}_m$ and $\widehat{a}_s$: the cutoff ability between L and M increases, while the cutoff between L and S decreases monotonically. This is equivalent to a continuous increase in S sector unit wages relative to L sector unit wages, and a continuous improvement in L sector unit wages compared to M sector unit wages. The bottom two panels show our model's predictions (solid lines) contrasted with the data (dashed lines) for our measures of interest. Not surprisingly, the model matches the 1950 employment shares (bottom left panel) and the 1950 relative average wages (bottom right panel) very well, as we targeted these measures. But the model also does extremely well in predicting the paths of the employment shares and relative average wages after 1950. In the data the employment share of the high skilled service sector increased by 17.38 percentage points, our model predicts a 14.95 percentage point increase. The employment share of manufacturing workers in the data fell by 21.36 percentage points, our model predicts a 22 percentage point contraction. Finally, the low-skilled service sector employment share increased by 3.98 percentage points in the data, whereas our model predicts a 7.1 percentage point increase.[26] In the data the relative average wage of low-skilled service workers compared to manufacturing workers increased by 15 per cent, while that of the high-skilled service workers increased by 28 per cent. In our model, as discussed in Section 3.3, the relative average wage changes are driven by changes in the relative unit wages and changes in the relative average sectoral abilities. As mentioned earlier, these two effects in general go in opposite directions, however the direct effect typically dominates the indirect effect of the changes in the optimal sector of work cutoffs. This is true in our baseline model as well: while low-skilled unit wages improved by 15 per cent relative to manufacturing unit wages, the manufacturing relative average ability improved by 4.5 per cent, leading to an overall 10 percent increase in the relative average low-skilled service sector wages. The high-skilled unit wage increased by 28 per cent relative to the manufacturing unit wage, but the high-skilled service sector average ability fell by almost 10 per cent compared to the manufacturing average ability, resulting in an overall 17 per cent increase in relative average high-skilled service sector wages. Thus, our baseline model predicts two thirds of the growth in low-skilled service sector wages, and 60 percent of the growth in relative average high-skilled service sector wages compared to manufacturing.

As we argued earlier, the differences in productivity growth between services and manufacturing are in general overestimated when not taking into account that due to the endogenous sorting of workers the average ability in each sector systematically changes. We therefore repeat our numerical exercise with the baseline calibration, but feeding in the average productivity growth rates adjusted for this bias. We assume that manufacturing grows at a slightly lower 2.18%, while services grow at a slightly higher 0.85% annual rate (bottom right numbers in Table 3). Figure 8 shows the dynamics of employment shares and relative average wages implied by the model in this case, which are very similar to the results of the baseline model. Since the relative annual productivity gain in manufacturing is lower once

[26]Since we are trying to match changes in shares, as long as some changes are understated, others are necessarily overstated.

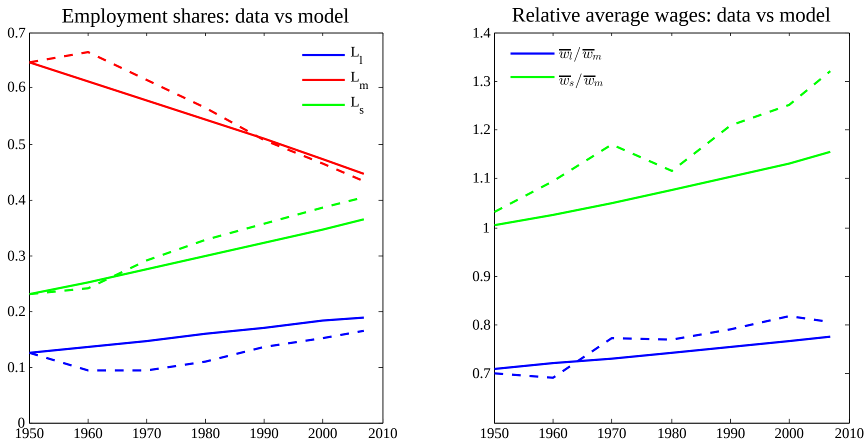

Figure 8: Transition of the model with selection-adjusted productivity growth rates

we correct for the changing selection of individuals, the equilibrium requires less labor to shift away from manufacturing. Our model now slightly under-predicts the contraction in manufacturing, and predicts a smaller expansion in both services, which brings the predictions closer to the data in terms of low-skilled services, and a bit further in terms of high-skilled services.

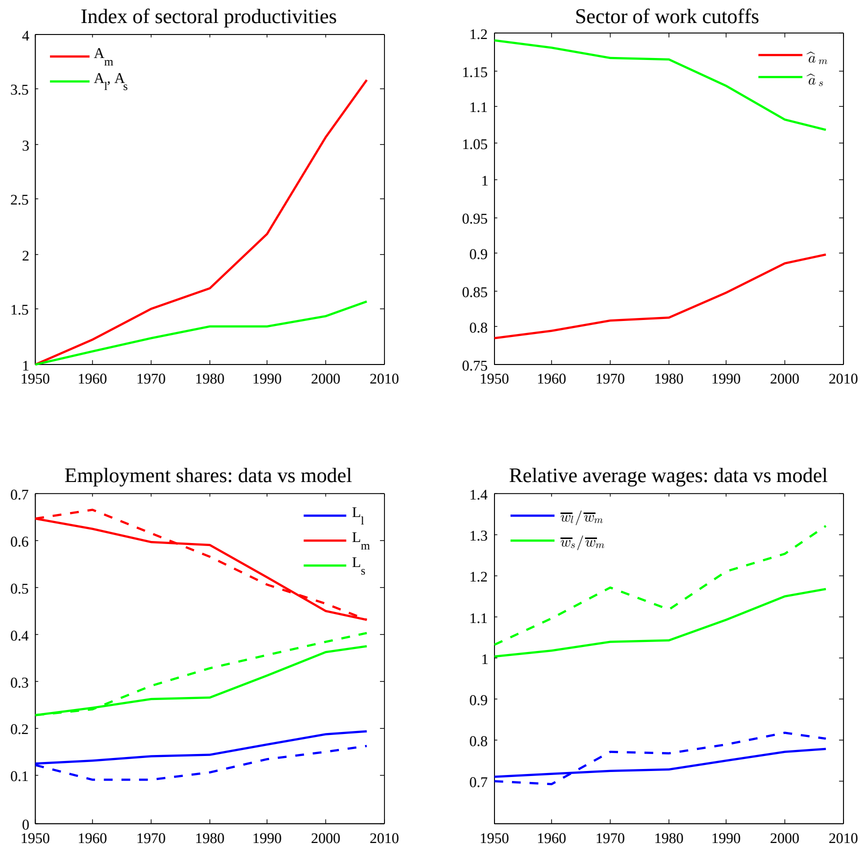

Figure 9: Transition of the model with decennial growth

The path of employment shares and relative average wages generated by the model are very smooth compared to the data. This is not surprising, as we assumed a constant annual growth rate of sectoral labor productivity between 1950 and 2007. However, Table 3 reveals that the growth rates have varied substantially over time. Figure 9 shows the simulated model contrasted with the data when feeding in the growth rates calculated for each period. The differential productivity gain of manufacturing across periods implies a less smooth change in employment shares and relative average wages across sectors, while the overall predicted changes are the same. Quantitatively, when feeding in the actual series of productivity growth (shown in the top left panel), the model matches the actual time path of these variables even better. In particular, in the data manufacturing productivity growth accelerated relative to services in 1980, and the model implies that from then on – in line with the data – the wage growth in high-skilled services compared to manufacturing increased, and that low-skilled service employment started to expand at a faster rate.

As discussed in Section 3.3, a model without capital intensity differences and with perfect capital mobility across sectors cannot match the data on relative average wages, sectoral employment shares and sectoral expenditure shares jointly. Since employment and wages are the focus of this paper, we chose these measures as targets in our calibration. Nonetheless, our model does quite a good job in matching the path of the relative value added share in manufacturing compared to combined services between 1950 and 2007. Figure 10 shows the relative value added in manufacturing compared to the value added in combined services in the model (solid line) and in the data (dashed line).[27] Even though the level of relative value added is not matched by the model, the overall decline is matched quite well: in the data it declined by 63.97 per cent, while in the model it declined by 64.89 per cent.

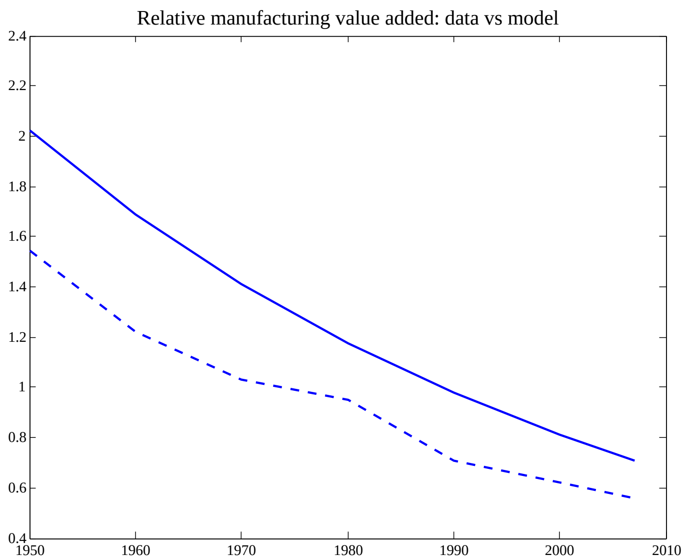

Figure 10: Transition of relative manufacturing value added

The graph shows the value added in manufacturing relative to combined services as predicted by the model (solid line) and in the data (dashed line) .

[27] We plot combined services, as the BEA data is not available for fine enough industry classification to calculate the value added in low- and high-skilled services from the data.

4.3 Robustness checks

In our baseline calibration we used $\varepsilon = 0.002$ for the elasticity of substitution between goods and services (measured in value-added terms) as estimated by Herrendorf, Rogerson, and Valentinyi (2013), which is at the lower end of estimates reported by Ngai and Pissarides (2008). To see whether our results are robust to higher, yet plausible, values of this parameter, we explore how our results change when using $\varepsilon = 0.02$ or $\varepsilon = 0.2$, naturally recalibrating the other parameters to match moments of the 1950 data.[28] Qualitatively the transition paths look exactly the same. A higher, but still low elasticity of substitution implies that the effective employment in low- and high-skilled services have to increase less, and the effective employment in manufacturing has to fall less in order to meet equilibrium demands. This in turn implies less adjustment in raw employment and in wages. Increasing the value of the elasticity of substitution brings the model's predictions for the time path of low-skilled service sector employment closer to the data, but worsens the fit of the other variables of interest. As can be seen in Table 8 in the Appendix, even with the least favorable calibration (high ε and labor productivity growth adjusted for selection), our model predicts 45 per cent of the increase in L and 39 per cent of the increase in S sector average wages relative to M. In terms of employment shares, the model's predictions are still within 20 per cent of the data in 2007.[29] Overall, the benchmark calibration with $\varepsilon = 0.002$ as estimated by Herrendorf, Rogerson, and Valentinyi (2013) seems to do best in replicating the data.

5 Conclusions

The literature on polarization of employment and wages has typically focused on occupations. We present a set of new empirical facts that suggest that in addition to reallocations between occupations within industries, also shifts between industries contribute to the polarization of labor markets. Moreover, we show that in terms of broadly defined industries, polarization was present as early as 1950-1960 and directly linked to the decline of manufacturing employment. Based on this evidence we propose a novel explanation, one based on structural change. A methodological contribution of our paper is that we develop a multi-sector model with heterogeneous labor in Roy-style fashion, the most parsimonious setup that yet allows heterogeneity in wages. An insight from our model is that unbalanced technological progress does not only lead to structural change, the reallocation of employment across sectors, but also affects sectoral average wages. We find that higher productivity growth in manufacturing than in services increases employment and wages in both the low-skilled and the high-skilled service sector, thus leading to the polarization of the labor market.

[28] For $\varepsilon = 0.02$ these are $\tau_l = 0.1744$ and $\tau_s = 0.5383$; for $\varepsilon = 0.2$ we obtain $\tau_l = 0.1671$ and $\tau_s = 0.4995$.

[29] While it predicts a 5.12 percentage point increase in L, a 15.03 percentage point fall in M, and a 9.92 percentage point increase in S employment share compared to the 3.98, 21.37 and 17.28 percentage point changes in the data.

References

ACEMOGLU, D., AND D. AUTOR (2011): "Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings," vol. 4, Part B of *Handbook of Labor Economics*, pp. 1043 – 1171. Elsevier.

ACEMOGLU, D., AND V. GUERRIERI (2008): "Capital Deepening and Nonbalanced Economic Growth," *Journal of Political Economy*, 116(3), pp. 467–498.

AUTOR, D. H., AND D. DORN (2013): "The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market," *American Economic Review*, 103(5), 1553–97.

AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2006): "The Polarization of the U.S. Labor Market," *The American Economic Review*, 96(2), pp. 189–194.

AUTOR, D. H., F. LEVY, AND R. J. MURNANE (2003): "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, 118(4), pp. 1279–1333.

BLINDER, A. S. (2009): "How Many US Jobs Might Be Offshorable?," *World Economics*, 10(2), 41–78.

BLINDER, A. S., AND A. B. KRUEGER (2013): "Alternative Measures of Offshorability: A Survey Approach," *Journal of Labor Economics*, 31(S1), S97 – S128.

BOPPART, T. (2014): "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences," *Econometrica*, forthcoming.

BUERA, F. J., AND J. P. KABOSKI (2012): "The Rise of the Service Economy," *American Economic Review*, 102(6), 2540–69.

CARNEIRO, P., AND S. LEE (2011): "Trends in Quality-Adjusted Skill Premia in the United States, 1960-2000," *American Economic Review*, 101(6), 2309–49.

CASELLI, F., AND W. J. I. COLEMAN (2001): "The U.S. Structural Transformation and Regional Convergence: A Reinterpretation," *Journal of Political Economy*, 109(3), pp. 584–616.

DORN, D. (2009): "Essays on Inequality, Spatial Interaction, and the Demand for Skills," Ph.D. thesis, University of St. Gallen, Dissertation no. 3613.

DUSTMANN, C., J. LUDSTECK, AND U. SCHÖNBERG (2009): "Revisiting the German Wage Structure," *The Quarterly Journal of Economics*, 124(2), 843–881.

GOOS, M., AND A. MANNING (2007): "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *The Review of Economics and Statistics*, 89(1), 118–133.

GOOS, M., A. MANNING, AND A. SALOMONS (2009): "Job Polarization in Europe," *The American Economic Review Papers and Proceedings*, 99(2), pp. 58–63.

——— (2014): "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," *American Economic Review*, 104(8), 2509–26.

GROSSMAN, G. M., AND E. ROSSI-HANSBERG (2008): "Trading Tasks: A Simple Theory of Offshoring," *American Economic Review*, 98(5), 1978–97.

HERRENDORF, B., R. ROGERSON, AND A. VALENTINYI (2013): "Two Perspectives on Preferences and Structural Transformation," *American Economic Review*, 103(7), 2752–89.

HERRENDORF, B., R. ROGERSON, AND . VALENTINYI (2014): "Growth and Structural Transformation," in *Handbook of Economic Growth*, vol. 2 of *Handbook of Economic Growth*, chap. 6, pp. 855–941. Elsevier.

KONGSAMUT, P., S. REBELO, AND D. XIE (2001): "Beyond Balanced Growth," *The Review of Economic Studies*, 68(4), pp. 869–882.

KUZNETS, S. (1957): "Quantitative Aspects of the Economic Growth of Nations: II. Industrial Distribution of National Product and Labor Force," *Economic Development and Cultural Change*, 5(4), pp. 1–111.

MADDISON, A. (1980): "Economic Growth and Structural Change in the Advanced Countries," in *Western Economies in Transition: Structural Change and Adjustment Policies in Industrial Countries*, ed. by I. Leveson, and J. W. Wheeler, pp. 41–65. London: Croom Helm.

MANNING, A. (2004): "We Can Work It Out: The Impact of Technological Change on the Demand for Low-Skill Workers," *Scottish Journal of Political Economy*, 51(5), 581–608.

MAZZOLARI, F., AND G. RAGUSA (2013): "Spillovers from high-skill consumption to low-skill labor markets," *The Review of Economics and Statistics*, 95(1), 74 – 86.

MEYER, P. B., AND A. M. OSBORNE (2005): "Proposed Category System for 1960-2000 Census Occupations," Working Papers 383, U.S. Bureau of Labor Statistics.

MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2014): "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 years," *The Review of Economics and Statistics*, 96(1), 60 – 77.

NGAI, L. R., AND B. PETRONGOLO (2014): "Gender Gaps and the Rise of the Service Economy," Discussion Paper 9970, Centre for Economic Policy Research.

NGAI, L. R., AND C. A. PISSARIDES (2007): "Structural Change in a Multisector Model of Growth," *The American Economic Review*, 97(1), pp. 429–443.

——— (2008): "Trends in hours and economic growth," *Review of Economic Dynamics*, 11(2), 239 – 256.

RUGGLES, S., J. T. ALEXANDER, K. GENADEK, R. GOEKEN, M. B. SCHROEDER, AND M. SOBEK (2010): "Integrated Public Use Microdata Series," Discussion paper, Minnesota Population Center, Minneapolis, MN, Version 5.0.

Spitz-Oener, A. (2006): "Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure," *Journal of Labor Economics*, 24(2), pp. 235–270.

Young, A. (2014): "Structural Transformation, the Mismeasurement of Productivity Growth, and the Cost Disease of Services," *American Economic Review*, 104(11), 3635–67.

Appendix

Data

We use data from the US Census of 1950, 1960, 1970, 1980, 1990, 2000 and the American Community Survey (ACS) of 2007, which we access from IPUMS-USA, provided by Ruggles, Alexander, Genadek, Goeken, Schroeder, and Sobek (2010). Following Acemoglu and Autor (2011) and Autor and Dorn (2013) we restrict the sample to individuals who were in the labor force and of age 16 to 64 in the year preceding the survey. We drop residents of institutional group quarters and unpaid family workers. We also drop respondents with missing earnings or hours worked data and those who work in agricultural occupations or industries or in the military. Our employment measure is the product of weeks worked times usual number of hours per week. We also compute hourly wages as earnings divided by the product of usual hours and weeks worked.[30]

Figure 11: Wage and employment polarization II.

Notes: Data and left and right panels as in Figure 1. The bottom two panels show the 10-year change in employment shares (calculated as hours supplied), and the top two panels show the 10-year change in log hourly real wages.

To construct the 30-year change graphs of Figure 1 and the 10-year change graphs of Figure 11 we

[30] Since in 1950 the Census did not include usual hours worked, we use hours worked last week instead. In 1960 and 1970 the Census asked only for an interval of hours and weeks worked last year; we use the midpoint of the interval given.

follow the methodology used in Autor, Katz, and Kearney (2006), Acemoglu and Autor (2011), and Autor and Dorn (2013), which requires a balanced panel of occupations. Dorn (2009) and Autor and Dorn (2013) provide a balanced panel of occupational classifications (‘*occ1990dd*’) over 1980-2008, which we use to construct a balanced panel over 1950-2007 by aggregating occupational codes as needed. This leaves us with 183 balanced occupational codes.

Categorization of occupations

Following the routinization literature, we classify occupations into three categories, which are used in Figure 3:

- Manual (low-skilled non-routine): housekeeping, cleaning, protective service, food prep and service, building, grounds cleaning, maintenance, personal appearance, recreation and hospitality, child care workers, personal care, service, healthcare support;
- Routine: construction trades, extractive, machine operators, assemblers, inspectors, mechanics and repairers, precision production, transportation and material moving occupations, sales, administrative support, sales, administrative support;
- Abstract (skilled non-routine): managers, management related, professional specialty, technicians and related support.

Categorization of industries

Based on our theory we classify the industries into three sectors, which are used in Figure 4:

- Low-skilled services: personal services, entertainment, business and repair services (except advertising and computer and data processing services), nursing and personal care facilities, child day care service, family child care homes, residential care facilities, social services not elsewhere classified, taxi, retail bakeries, eating and drinking places;
- Manufacturing: mining, construction, manufacturing, transport and public utilities, wholesale trade, retail trade;
- High-skilled services: communications, finance, insurance and real estate, theaters and motion pictures, professional and related services, public administration, advertising, computer and data processing services.

Table 4 summarizes the descriptive statistics for sectoral employment:

Table 4: Descriptive statistics by industry

	low-skilled services	manufacturing	high-skilled services
Highschool Dropout	23.79%	23.84%	6.60%
Highschool Graduate	34.95%	38.04%	22.29%
Some College	27.90%	24.41%	29.09%
College Degree	10.35%	10.81%	24.67%
Postgraduate	3.01%	2.90%	17.35%
Mean Years of Education	12.21	12.17	14.31
Female Share	54.12%	27.29%	56.04%
Foreign-Born Share	16.18%	10.20%	9.03%

Occupation and sector premia

In Figures 3 and 4 as well as in our quantitative exercise we focused on relative average residual wages. We obtain these by regressing log hourly wages on dummies for industries and on a set of controls, comprising of a polynomial in potential experience (defined as age - years of schooling - 6), dummies for gender, race, and born abroad. Tables 5 and 6 show the regression results. Since we omit the dummy for manufacturing, the implied relative wage of a low-skilled (high-skilled) service worker is given by the exponential of the estimated coefficient on the low-skilled (high-skilled) service sector dummy. The regression specification to compute residual occupational wages is analogue, with the industry dummies replaced by occupation dummies; we omit the dummy for routine occupations, such that relative wages compared to routine occupations are given by the exponential of the sectoral dummies of low- and high-skilled services.

Table 5: Regression of log hourly wages: industry effects

Year	1950	1960	1970	1980	1990	2000	2007
low-sk. serv.	-0.36***	-0.37***	-0.26***	-0.26***	-0.24***	-0.20***	-0.22***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
high-sk. serv.	0.03***	0.09***	0.16***	0.11***	0.19***	0.22***	0.28***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Controls	yes	yes	yes	yes	yes	yes	yes
Observations	113635	459564	579290	958318	1094458	1235282	1308885
R^2	0.18	0.19	0.16	0.19	0.21	0.19	0.22

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 6: Regression of log hourly wages: occupation effects

Year	1950	1960	1970	1980	1990	2000	2007
manual	-0.35***	-0.37***	-0.27***	-0.24***	-0.22***	-0.18***	-0.12***
	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
abstract	0.21***	0.30***	0.35***	0.32***	0.40***	0.45***	0.50***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Controls	yes	yes	yes	yes	yes	yes	yes
Observations	113635	459564	579290	958318	1094458	1235282	1308885
R^2	0.20	0.23	0.20	0.23	0.26	0.25	0.27
Year	1950	1960	1970	1980	1990	2000	2007

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

In the text we showed the relative average occupational wages, and the coefficient on the sectoral dummies from wage regressions. In Figure 12 we show the reverse: the sectoral relative average wages compared to manufacturing, and the coefficients on occupational dummies from a wage regression. The patterns are left unchanged.

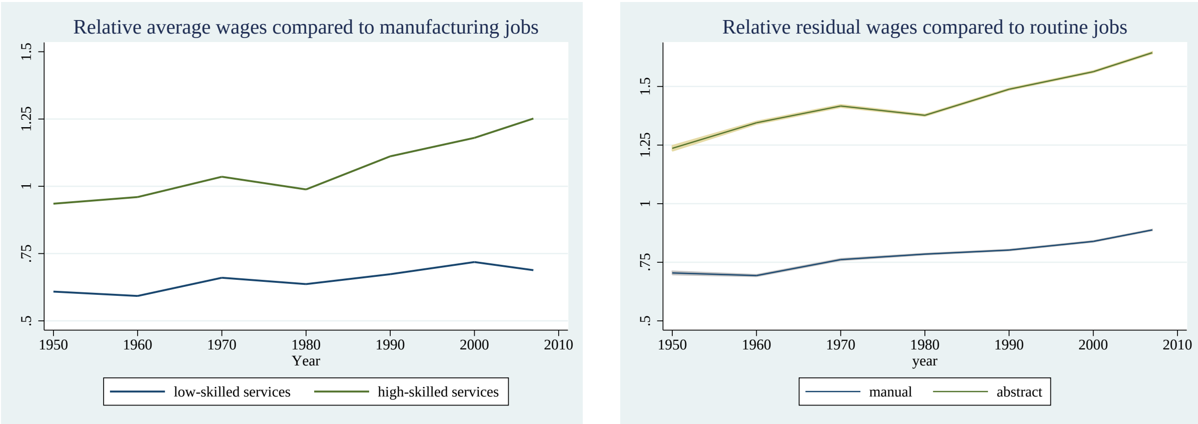

Figure 12: Wage polarization for sectors and occupations

Notes: Same data and classification as in Figure 3 and Figure 4. The left panel shows the relative average wages of high-skilled and low-skilled service workers compared to manufacturing workers. The right panel shows the occupation premium for abstract and manual workers compared to routine workers, and their 95 % confidence intervals, as estimated in Table 6.

The role of gender and age composition changes

Figure 13 demonstrates that the sectoral employment share changes are not driven by changes in the age, gender, race composition of the labor force. The counterfactual industry employment shares are generated by fixing the sectoral employment share of each age-gender-race cell at its 1950 level, and allowing the employment shares of the cells to change. While it can be seen that the counterfactual employment shares (the dashed lines) qualitatively move in the same direction as the actual employment shares (the solid lines), in terms of magnitude the counterfactual employment shares move much less.

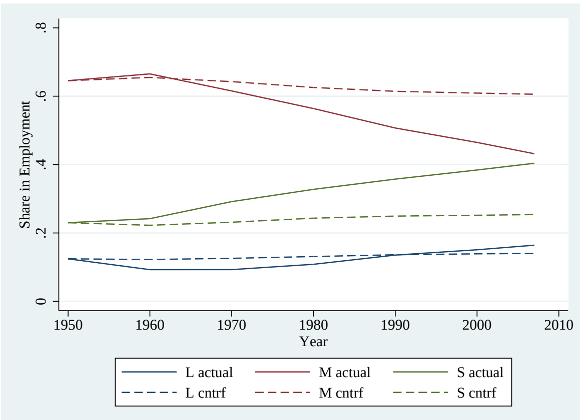

Figure 13: Counterfactual exercise: only changes in the gender-age composition of the labor force

Notes: Employment shares (in terms of hours) are calculated from the same data as in Figure 4. The actual data is shown as solid lines, while the dashed line show how the employment shares of industries would have evolved if only the relative size of gender-age cells in the labor force had changed over time.

The role of industry shifts in occupational employment shares

In the table below we show the decomposition of employment share and average wage changes for the entire 1950-2007 period, and also for two shorter time spans: 1950-1980, and 1980-2007.

Table 7: Decomposition of the changes in occupational employment shares

	3 occupations, 3 sectors			**10 occupations, 11 industries**		
	1950 - 1980	1980 - 2007	1950 - 2007	1950 - 1980	1980 - 2007	1950 - 2007
Manual						
Total Δ	-2.11	5.09	2.98	-2.08	5.21	3.12
Between Δ	0.50	3.49	4.39	-0.40	4.09	3.98
Within Δ	-2.61	1.59	-1.41	-1.68	1.12	-0.85
Routine						
Total Δ	-7.59	-12.20	-19.79	-10.94	-14.68	-25.80
Between Δ	-3.41	-6.51	-10.46	-4.23	-6.84	-12.38
Within Δ	-4.18	-5.69	-9.33	-6.70	-7.84	-13.42
Abstract						
Total Δ	9.70	7.11	16.81	10.65	8.93	19.79
Between Δ	2.90	3.02	6.07	4.17	3.76	8.94
Within Δ	6.79	4.09	10.74	6.47	5.17	10.84
Average						
Total Δ	-9.37	-11.79	-21.16	-11.05	-9.00	-20.04
Between Δ	-4.61	-6.67	-12.05	-4.07	-3.82	-8.83
Within Δ	-4.76	-5.12	-9.10	-6.98	-5.18	-11.21

Notes: Same data as in Figure 1. For each occupational category, the first row presents the change in the share of employment (in terms of hours worked), the second the between industry component, and the third the within industry component for the given time interval. The first three columns use 3 occupations and 3 sectors, the second three columns use 10 occupations and 11 industries.

An alternative way to asses the importance of the employment reallocations between industries for the shifts in the broad occupation categories, we conduct the following counterfactual exercise: we fix the industry shares in employment (in terms of hours worked) at their 1950 levels and let the within-industry share of occupations follow their actual path, and compute how the occupational shares would have evolved in the absence of between industry shifts. Figure 14 shows the resulting time series (dashed) and the actual data (solid). The exercise shows that if there were only within-industry shifts qualitatively the employment of the occupation categories would have evolved as in the actual data, but the exercise also points out that quantitatively they cannot explain all of the reallocation. We therefore conclude that also between industry shifts account for the polarization of occupational employment.

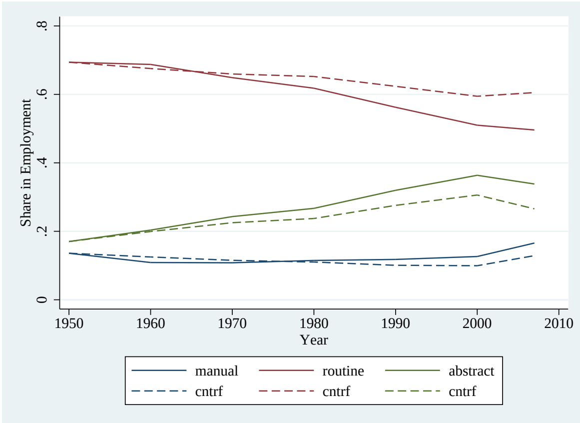

Figure 14: Counterfactual exercise: only-within industry shift of occupations

Notes: Employment shares (in terms of hours) are calculated from the same data as in Figure 12. The actual data is shown as solid lines, while the dashed line show how the occupational employment shares would have evolved in the absence of reallocations across industries.

Model

Proof of Proposition 1. Starting from:

$$\frac{L_l(\widehat{a}_m,\widehat{a}_s)}{N_m(\widehat{a}_m,\widehat{a}_s)}\widehat{a}_m^{\varepsilon}=\left(\frac{A_m}{A_l}\right)^{1-\varepsilon}\left(\frac{\theta_m}{\theta_l}\right)^{-\varepsilon}$$
$$\frac{N_s(\widehat{a}_m,\widehat{a}_s)}{N_m(\widehat{a}_m,\widehat{a}_s)}\left(\frac{\widehat{a}_m}{\widehat{a}_s}\right)^{\varepsilon}=\left(\frac{A_m}{A_s}\right)^{1-\varepsilon}\left(\frac{\theta_m}{\theta_s}\right)^{-\varepsilon}$$

A change in productivities triggers changes in the equilibrium cutoffs, $\widehat{a}_m$ and $\widehat{a}_s$, in such a way that the above conditions remain satisfied. Total differentiation then implies:

$$\varepsilon\frac{d\widehat{a}_m}{\widehat{a}_m}+\frac{dL_l}{L_l}-\frac{dN_m}{N_m}=\frac{d\frac{A_m}{A_l}}{\frac{A_m}{A_l}} \tag{21}$$
$$\varepsilon\left(\frac{d\widehat{a}_m}{\widehat{a}_m}-\frac{d\widehat{a}_s}{\widehat{a}_s}\right)+\frac{dN_s}{N_s}-\frac{dN_m}{N_m}=\frac{d\frac{A_m}{A_s}}{\frac{A_m}{A_s}} \tag{22}$$

Applying the Leibnitz rule to the expressions for $L_l(\widehat{a}_m,\widehat{a}_s), N_m(\widehat{a}_m,\widehat{a}_s)$ and $N_s(\widehat{a}_m,\widehat{a}_s)$, we get the following expressions for the change in the effective and raw labor supplies as a function of the change in $\widehat{a}_m$ and in $\widehat{a}_s$ is

$$dL_l(\widehat{a}_m,\widehat{a}_s)=\frac{\partial n_l}{\partial\widehat{a}_m}d\widehat{a}_m+\frac{\partial n_l}{\partial\widehat{a}_s}d\widehat{a}_s=\underbrace{\int_0^{\widehat{a}_s}f(\widehat{a}_m,a_s)da_s}_{\equiv C_1}\cdot d\widehat{a}_m+\underbrace{\int_0^{\widehat{a}_m}f(a_m,\widehat{a}_s)da_m}_{\equiv C_2}\cdot d\widehat{a}_s \tag{23}$$
$$dN_m(\widehat{a}_m,\widehat{a}_s)=-\underbrace{\int_0^{\widehat{a}_s}f(\widehat{a}_m,a_s)da_s}_{=C_1}\cdot\widehat{a}_m d\widehat{a}_m-\underbrace{\int_{\widehat{a}_m}^{\infty}a_m^2f(a_m,\frac{\widehat{a}_s}{\widehat{a}_m}a_m)da_m}_{\equiv C_3}\cdot\frac{\widehat{a}_s}{\widehat{a}_m}\left(\frac{d\widehat{a}_m}{\widehat{a}_m}-\frac{d\widehat{a}_s}{\widehat{a}_s}\right) \tag{24}$$

$$dN_s(\widehat{a}_m,\widehat{a}_s) = -\underbrace{\int_0^{\widehat{a}_m} f(a_m,\widehat{a}_s)da_m}_{=C_2} \cdot \widehat{a}_s d\widehat{a}_s + \underbrace{\int_{\widehat{a}_s}^{\infty} a_s^2 f(\frac{\widehat{a}_m}{\widehat{a}_s}a_s, a_s)da_s}_{\equiv C_4} \cdot \frac{\widehat{a}_m}{\widehat{a}_s}\left(\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s}\right) \tag{25}$$

similarly

$$dL_m(\widehat{a}_m,\widehat{a}_s) = -\underbrace{\int_0^{\widehat{a}_s} f(\widehat{a}_m,a_s)da_s}_{=C_1} \cdot d\widehat{a}_m - \underbrace{\int_{\widehat{a}_m}^{\infty} a_m f(a_m,\frac{\widehat{a}_s}{\widehat{a}_m}a_m)da_m}_{\equiv C_5} \cdot \frac{\widehat{a}_s}{\widehat{a}_m}\left(\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s}\right) \tag{26}$$

$$dL_s(\widehat{a}_m,\widehat{a}_s) = -\underbrace{\int_0^{\widehat{a}_m} f(a_m,\widehat{a}_s)da_m}_{=C_2} \cdot d\widehat{a}_s + \underbrace{\int_{\widehat{a}_s}^{\infty} a_s f(\frac{\widehat{a}_m}{\widehat{a}_s}a_s, a_s)da_s}_{\equiv C_6} \cdot \frac{\widehat{a}_m}{\widehat{a}_s}\left(\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s}\right) \tag{27}$$

Plugging these into (21) and (22) and re-arranging we get:

$$\frac{d\widehat{a}_m}{\widehat{a}_m}\underbrace{\left[\varepsilon + \frac{C_1\widehat{a}_m}{L_l} + \frac{C_1(\widehat{a}_m)^2}{N_m} + \frac{C_3\frac{\widehat{a}_s}{\widehat{a}_m}}{N_m}\right]}_{\equiv B_1>0} + \frac{d\widehat{a}_s}{\widehat{a}_s}\underbrace{\left[\frac{C_2\widehat{a}_s}{L_l} - \frac{C_3\frac{\widehat{a}_s}{\widehat{a}_m}}{N_m}\right]}_{\equiv B_2} = \underbrace{\frac{d\frac{A_m}{A_l}}{\frac{A_m}{A_l}}}_{\equiv D_1}$$

$$\frac{d\widehat{a}_m}{\widehat{a}_m}\underbrace{\left[\varepsilon + \frac{C_4\frac{\widehat{a}_m}{\widehat{a}_s}}{N_s} + \frac{C_1(\widehat{a}_m)^2}{N_m} + \frac{C_3\frac{\widehat{a}_s}{\widehat{a}_m}}{n_m}\right]}_{\equiv B_3>0} - \frac{d\widehat{a}_s}{\widehat{a}_s}\underbrace{\left[\varepsilon + \frac{C_2(\widehat{a}_s)^2}{N_s} + \frac{C_4\frac{\widehat{a}_m}{\widehat{a}_s}}{N_s} + \frac{C_3\frac{\widehat{a}_s}{\widehat{a}_m}}{N_m}\right]}_{\equiv B_4>0} = \underbrace{\frac{d\frac{A_m}{A_s}}{\frac{A_m}{A_s}}}_{\equiv D_2}$$

This leads to

$$\frac{d\widehat{a}_s}{\widehat{a}_s} = \frac{B_3D_1 - B_1D_2}{B_3B_2 + B_1B_4}$$
$$\frac{d\widehat{a}_m}{\widehat{a}_m} = \frac{D_2B_2 + B_4D_1}{B_3B_2 + B_1B_4}$$

where $B_3B_2 + B_1B_4 > 0$ always holds. Hence to determine the response in $\widehat{a}_m$ and in $\widehat{a}_s$, we only need to consider the sign of the numerator. If $D_1 = D_2 > 0$, i.e. the growth rate of A_l is equal to the growth rate of A_s, and lower than the growth rate of A_m, then the following expressions can be obtained:

$$\frac{d\widehat{a}_s}{\widehat{a}_s} = \frac{D}{B_3B_2 + B_1B_4}(B_3 - B_1) = \frac{D}{B_3B_2 + B_1B_4}\left(\frac{C_4\frac{\widehat{a}_m}{\widehat{a}_s}}{N_s} - \frac{C_1\widehat{a}_m}{L_l}\right), \tag{28}$$

$$\frac{d\widehat{a}_m}{\widehat{a}_m} = \frac{D}{B_3B_2 + B_1B_4}(B_2 + B_4) = \frac{D}{B_3B_2 + B_1B_4}\left(\varepsilon + \frac{C_2(\widehat{a}_s)^2}{N_s} + \frac{C_4\frac{\widehat{a}_m}{\widehat{a}_s}}{N_s} + \frac{C_2\widehat{a}_s}{L_l}\right) > 0. \tag{29}$$

As this shows, $\frac{d\widehat{a}_m}{\widehat{a}_m} > 0$. The sign of $\frac{d\widehat{a}_s}{\widehat{a}_s}$ is ambiguous in general, but it is straightforward that $\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s} > 0$:

$$\left(\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s}\right) = \frac{D}{B_3B_2 + B_1B_4}\left(\varepsilon + \frac{C_2\widehat{a}_s}{L_l} + \frac{C_2(\widehat{a}_s)^2}{N_s} + \frac{C_1\widehat{a}_m}{L_l}\right) > 0$$

This together with (24) and (26) imply that N_m and L_m always decrease. These changes are:

$$dN_m(\widehat{a}_m,\widehat{a}_s) = -\left(C_1(\widehat{a}_m)^2\frac{d\widehat{a}_m}{\widehat{a}_m} + C_3\frac{\widehat{a}_s}{\widehat{a}_m}\left(\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s}\right)\right) < 0$$
$$dL_m(\widehat{a}_m,\widehat{a}_s) = -\left(C_1\widehat{a}_m\frac{d\widehat{a}_m}{\widehat{a}_m} + C_5\frac{\widehat{a}_s}{\widehat{a}_m}\left(\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s}\right)\right) < 0$$

By plugging in (28) and (29) into (23) we can show that employment in sector L increases:

$$\begin{aligned} dL_l(\widehat{a}_m,\widehat{a}_s) &= C_1\widehat{a}_m\frac{d\widehat{a}_m}{\widehat{a}_m} + C_2\widehat{a}_s\frac{d\widehat{a}_s}{\widehat{a}_s} \\ &= \frac{D}{B_3B_2+B_1B_4}\left[C_1\widehat{a}_m\left(\varepsilon + \frac{C_2(\widehat{a}_s)^2}{N_s} + \frac{C_4\frac{\widehat{a}_m}{\widehat{a}_s}}{N_s}\right) + C_2\widehat{a}_s\frac{C_4\frac{\widehat{a}_m}{\widehat{a}_s}}{N_s}\right] > 0. \end{aligned}$$

By plugging in (28) and (29) into (25) we can show that effective employment in sector S increases:

$$dN_s(\widehat{a}_m,\widehat{a}_s) = \frac{D}{B_3B_2+B_1B_4}\left[C_2\widehat{a}_s^2\frac{C_1\widehat{a}_m}{L_l} + C_4\frac{\widehat{a}_m}{\widehat{a}_s}\left(\varepsilon + \frac{C_2\widehat{a}_s}{L_l} + \frac{C_1\widehat{a}_m}{L_l}\right)\right] > 0.$$

□

Quantitative results

Labor productivity calculation

To calculate labor productivity growth for our classification of industries, similarly to Ngai and Petrongolo (2014), we divide the sectoral value added by the sectoral employment. To calculate our sectoral value added data we take industry level value added data from Herrendorf, Rogerson, and Valentinyi (2013), which are based on the BEA data but are corrected for the use of any industry's value added as investment. To ensure consistency with the sector classification in the value-added data, we use data on employment from the BEA, which for the time span that we need (1950-2007) in only available in terms of number of employees. As it is not possible to distinguish between low- and high-skilled services in this data, we aggregate up both the value-added data from Herrendorf et al (2013) and the BEA employment data to the manufacturing sector and a combined service sector given our classification. Then we compute labor productivity in manufacturing and in services by dividing the value-added by the employment for each year.

To correct for the ability bias that arises when workers self-select into sectors, we make use of the functional form of the ability distribution which we have calibrated. Given this distribution we can calculate the cutoff abilities $\widehat{a}_m$ and $\widehat{a}_s$ needed to match the employment shares calculated from the Census and ACS data. Given the cutoffs and the distribution of abilities, we can calculate the average ability in each sector for these years (1950, 1960, 1970, 1980, 1990, 2000, 2007). We then generate effective employment from these average abilities and the raw employment data from the BEA. The adjusted labor productivity is then just the sectoral value added divided by the effective employment.

Employment and wage dynamics with decennial and adjusted labor productivity

Figure 15 shows the model implied wage and employment dynamics when feeding in the growth rates computed for each model period adjusted for self-selection (the last two columns of Table 3).

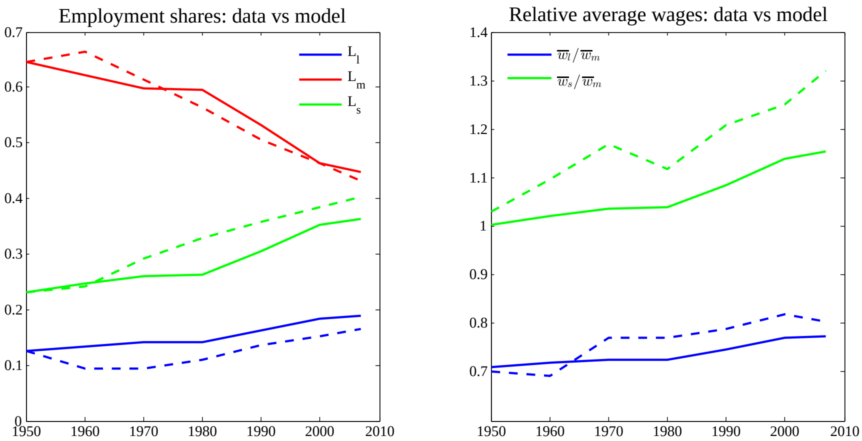

Figure 15: Transition of the model with selection-adjusted decennial productivity growth rates

Robustness

Table 8: Robustness checks: different model specifications vs the data

		employment share Δ			relative average wages Δ	
ε	productivity	L	M	S	L to M	S to M
0.002	raw	7.11	-22.07	14.96	10.08	16.99
	adjusted	6.45	-19.88	13.43	9.02	15.04
0.02	raw	6.98	-21.55	14.57	9.85	16.50
	adjusted	6.34	-19.42	13.08	8.82	14.61
0.2	raw	5.66	-16.70	11.03	7.64	12.15
	adjusted	5.12	-15.03	9.92	6.85	10.84
	data	**3.98**	**-21.37**	**17.38**	**15.00**	**27.97**

Notes: The third to fifth columns shows the percentage point change in employment shares, while the last two columns show the percentage change in relative average wages. The first six rows contain the implied change for different elasticities of substitutions (ε) and for raw and adjusted labor productivity growth rates (from Table 3). The last row contains the change in these same measures between 1950 and 2007 in the data.