

Job Polarization and Structural Change*

Zsófia L. Bárány[†]

Christian Siegel[‡]

27 June 2016

Abstract

We document that job polarization – contrary to the consensus – has started as early as the 1950s in the US: middle-wage workers have been losing both in terms of employment and average wage growth compared to low- and high-wage workers. Given that polarization is a long-run phenomenon and closely linked to the shift from manufacturing to services, we propose a structural change driven explanation, where we explicitly model the sectoral choice of workers. Our simple model does remarkably well not only in matching the evolution of sectoral employment, but also of relative wages over the past fifty years.

JEL codes: E24, J22, O41

Keywords: Job Polarization, Structural Change, Roy model

*We wish to thank Francesco Caselli, Nicolas Coeurdacier, David Dorn, Georg Duernecker, Georg Graetz, Emeric Henry, Per Krusell, L. Rachel Ngai, Rish Singhania, 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, the German Economic Association Annual Meeting 2014, the 23rd CEPR European Summer Symposium in International Macroeconomics 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 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-wage occupations, also entails the relative growth of wages and employment of low-wage occupations compared to middle-wage 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-wage (routine) occupations, whereas they complement the high-skilled and high-wage (abstract) occupations (Autor, Levy, and Murnane (2003), Autor, Katz, and Kearney (2006), Autor and Dorn (2013), Feng and Graetz (2014), Goos, Manning, and Salomons (2014), Michaels, Natraj, and Van Reenen (2014)).

The contribution of our paper is twofold. First, we document a set of facts which raise flags that routinization, although certainly playing a role from the 1980s onwards, might not be the only driving force 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 data for the period 1950-2007 reveals some novel facts.² First, we document that polarization defined over occupational categories both in terms of employment and wages has been present in the US since the 1950s, which is long before ICT could have played a role. Second, we show that at least since the 1960s the same patterns for both employment and wages are discernible in terms of three broad sectors: low-skilled services, manufacturing and high-skilled services. Additionally, we confirm previous findings that a significant part of the observed occupational employment share changes are driven by sectoral employment shifts; thus understanding the sectoral labor market trends is important even for the occupational trends.

Based on these facts, we propose a structural change driven explanation for these

¹See 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, 2014) for Europe.

²Analyzing 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.

sectoral labor market trends. 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 complements, a change in relative productivities increases labor demand in the relatively slow growing sectors, and wages in these sectors have 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.³ We break services into two as they comprise of many different subsets, e.g. dry cleaners vs. banking, which seem hardly to be perfect substitutes in consumption, as would be implied by having a single service consumption in households' preferences. In our model, we therefore treat low- and high-end services as being just as substitutable with each other as they are with manufacturing goods.

A change in relative productivities does not only affect relative supply, but through prices it also affects relative demand. Given that goods and the two types of services are complements, as relative labor productivity in manufacturing increases, labor has to reallocate from manufacturing to both service sectors. To attract more workers into both low- and high-skilled services, their wages have to improve relative to manufacturing. Since in the data we see that manufacturing jobs tend to be in the middle of the wage distribution, this mechanism leads to a pattern of polarization in terms of sectors, which is driven by the interaction of supply and demand for sectoral output.

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 more than 70 per cent of the relative average wage gain of high- and low-skilled services compared to manufacturing, and around three quarters of the change in employment shares.

³Buera 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.

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. Our model highlights a particular relationship between structural change and job polarization. Structural change leads to the sector with the highest productivity growth to shrink in terms of employment, and to experience lower wage growth than the other sectors of the economy. If this sector is the most intensive in the occupations which are “hollowing out”, then structural change leads to polarization across occupations.

The structural change literature has documented for several countries that as income increases employment shifts away from agriculture and from manufacturing towards services, and 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 1950/60s 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 the employment patterns are mimicked by the path of relative average wages. The economic mechanisms put forward in the literature for 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 sectors (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’

⁴Acemoglu and Autor (2011) and Goos et al. (2014) look at the contribution of between-industry shifts to the polarization of occupational employment, but do not analyze the effect of structural change on the polarization of the labor market.

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. The change in relative sectoral labor productivity can be driven by differential sectoral TFP changes or by capital accumulation and different sectoral capital intensities.⁶ We stay agnostic about the origin of the differential labor productivity growth across sectors, and as Goos et al. (2014) point out it is possible that part of this since the 1980s or 1990s is driven by different routine intensities and ICT.

Ours is not the first paper to consider sectoral choice in a model of structural change. The setup of Matsuyama (1991) is similar, where agents have different efficiencies across sectors, but focuses on the theoretical possibility of multiplicity of stationary steady states. Caselli and Coleman (2001) study the role of falling costs of education in the structural shift from agriculture to manufacturing, and they derive predictions about the relative wages in the farm and non-farm sector. Focusing on cross-country differences, Lagakos and Waugh (2013) show that self-selection can account for gaps in productivity and wages between agriculture and non-agriculture. Buera and Kaboski (2012) analyze the relation between the increasing value added share of the service sector and the increasing skill premium, without exploring their model's implications for sectoral employment or wages, whereas this is the focus of our paper.

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. As mentioned before, the leading explanation is routinization linked to ICT. 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

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

⁶For example if ICT is used more intensively in the manufacturing sector, and ICT becomes cheaper, then this would show up as an increase in the relative productivity of manufacturing workers.

observed earlier.⁷ Another 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)). We do not incorporate such a mechanism in our model, as we strive for the most parsimonious setup featuring structural change, which does a good job in replicating the basic sectoral labor market facts since the 1960s.

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.

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 as well, 3) a significant part of employment polarization in terms of occupations is driven by employment shifts across industries. The focus of our quantitative model is fact 2). We document fact 1) as most of the literature documents polarization in terms of occupations. Fact 3) has been documented in the literature (Acemoglu and Autor (2011), Goos et al. (2014)), here we confirm it for our classification, data, and time horizon. We show fact 3) to convince the reader whose main interest is in occupations, that for the full picture one needs to consider industries as well. In what follows we document each of these facts in detail.

⁷Another 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)). Just as for the routinization hypothesis, this mechanism could have explanatory power from the 1980s onwards.

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 two occupational classifications: we start from the finest balanced occupational codes possible, and then go to ten broad occupational categories.

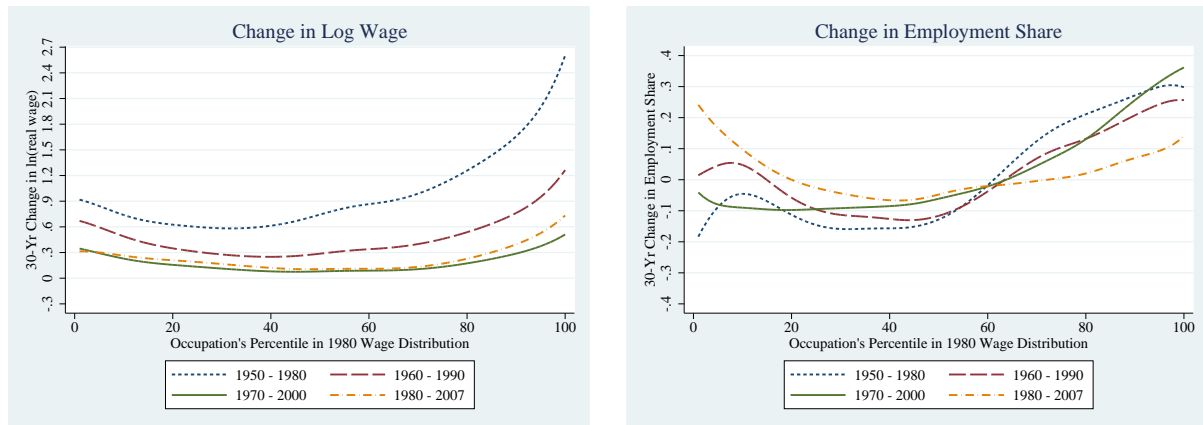


Figure 1: Smoothed changes in wages and employment

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 horizontal axis contains occupational skill percentiles based on their 1980 mean wages (see appendix for details). In the left panel the vertical axis shows for each occupational skill percentile the 30-year change in log hourly real wages, whereas in the right panel it shows the 30-year change in employment shares (calculated as hours supplied).

Following the methodology used in Autor et al. (2006), Acemoglu and Autor (2011), and Autor and Dorn (2013), we plot the smoothed changes in log real wages and employment shares for occupational percentiles, where occupations are ranked according to their 1980 mean hourly wages.⁸ The novelty in these graphs is that we show these patterns going back until 1950, rather than focusing only on the post-1980 period.⁹ In both graphs, each of the four curves represent changes which occurred over a different 30-year period. The left panel in Figure 1 shows that there has been polarization in terms of real wages in all 30-year periods, since the real wage change is

⁸We split occupations into 100 groups, each representing 1 percent of employment in 1980. We smooth changes in log real hourly wages and employment shares with a locally weighted regression using a bandwidth of 0.8.

⁹For comparability with the literature, we rank occupations based on their mean hourly wage in 1980. However, given that we look at a longer horizon than most of the literature, we also plot these changes against a different ranking of occupations, one based on the 1950 mean hourly real wages. The patterns look the same, see Figure 10 and the discussion in the appendix.

larger for low- and for high-ranked occupations than it is for middle-ranked occupations. 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 right panel shows the smoothed employment share changes. The picture shows that employment did not move monotonically towards higher wage occupations, instead it seems that middle-earning occupations lost the most in terms of employment. Thus employment polarization is present in the sense that the employment share in low- and high-wage occupations increased more (or decreased less) than in middle-wage occupations. Polarization in terms of employment is most pronounced in the last 30 years (1980-2007), but it seems to be present even in the earlier periods.¹⁰

We focus on 30-year windows for two reasons. First, most of the literature documents polarization over periods longer than two decades. Second, we link polarization to structural change, which is a long-run phenomenon. One concern with showing 30-year windows is that they stay silent on the exact timing of when polarization started. To address this, we show a decade-by-decade version of Figure 1 in the appendix. Figure 11 shows that these patterns do not necessarily hold for a decade-by-decade analysis, neither in the earlier nor in the later part of the sample. 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.¹¹ However, between 1960 and 1970 there is clear evidence of polarization, therefore the early polarization patterns in the 30-year windows are not solely driven by changes around 1980.

A set of balanced occupational categories is needed to generate Figure 1. Meyer and Osborne (2005) develop a set of harmonized occupational codes over the 1950 to 2000 Census and the ACS data. Dorn (2009) aggregates Meyer and Osborne (2005) to achieve the finest possible yet balanced set of categories from 1980 onward. We base our categories on Meyer and Osborne (2005) and Dorn (2009) to similarly achieve the finest possible balanced set of occupations from 1950 onward. One concern with this

¹⁰When ranking occupations in terms of the 1950, i.e. the initial, wage distribution, as is commonly done in this literature, employment polarization is much more noticeable since the 1960s, see Figure 10 in the appendix.

¹¹The fact that over 10-year windows the polarization patterns do not hold is in line with the evidence in Acemoglu and Autor (2011), see their Figure 10.



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 (left panels) and in its employment share (right panels) against its median log wage in 1980. The size of the circles represents the occupation's employment share in the initial year. The top panels show the changes between 1950 and 1980, whereas the bottom panels show the changes between 1980-2007. The 10 occupational categories 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.

approach is that despite the efforts of these authors, their harmonized and/or balanced categories are not truly comparable across Census years, and the reader might worry, that looking at a longer horizon, as we do, only exacerbates this problem. However, the biggest change in occupational (and industry) classification occurred with the implementation of the Standard Occupational Classification (SOC) based occupation (and the North American Industrial Classification System (NAICS) based industry) codes in the 2000 Census, when the hierarchical structure of occupational codes (and industries) was drastically modified. While in previous Census years certain smaller occupation (and industry) categories disappeared or entered, the main structure of occupations (and industries) remained the same.¹²

¹²See the documentation of the Integrated Public Use Microdata Series (IPUMS) Census data ¹³.

To minimize the comparability issues of fine occupational categories arising from the reclassifications across Census years, we aggregate up these fine categories to a coarser set of occupations. For these coarser categories there is less of a concern about the consistency over time.

Similarly to Goos and Manning (2007) and Goos et al. (2009), Figure 2 shows for 10 occupational groups the change in each occupation's median log wage (left panels) and in its employment share (right panels) against its median log wage in 1980. The panels on the top show the change between 1950-1980 and the panels on the bottom 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. The U-shape displayed in the plots confirms polarization of real wages and employment, both for the 1950-1980 and the 1980-2007 period.¹⁴

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. Our classification for the manufacturing sector includes also mining and construction, as is common in the structural change literature (e.g. Herrendorf, Rogerson, and Valentinyi (2013)). As mentioned in the introduction, we split the remaining (service) industries into two categories, where within sectors the industries should be close substitutes, whereas across sectors they should be complements. Our classification is also guided by production side considerations; in low-end services workers have much less education and earn much lower wages than in high-end. For this reason we refer to them as low-skilled and high-skilled services.¹⁵ As a result of the combined production and consumption side considerations we classify as low-end services the following industries: personal services, entertainment, low-skilled transport, low-skilled business

¹⁴In Figure 12 in the appendix, we document polarization in terms of occupations in an even coarser classification. Following Acemoglu and Autor (2011) we classify occupation groups into three categories: manual, routine, and abstract. Again, we find that 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.

¹⁵See Figure 3, as well as Figure 13 and Table 5 in the appendix.

and repair services, retail trade, and wholesale trade. High-end services comprise of professional and related services, finance, insurance and real estate, communications, high-skilled business services, utilities, high-skilled transport, and public administration.

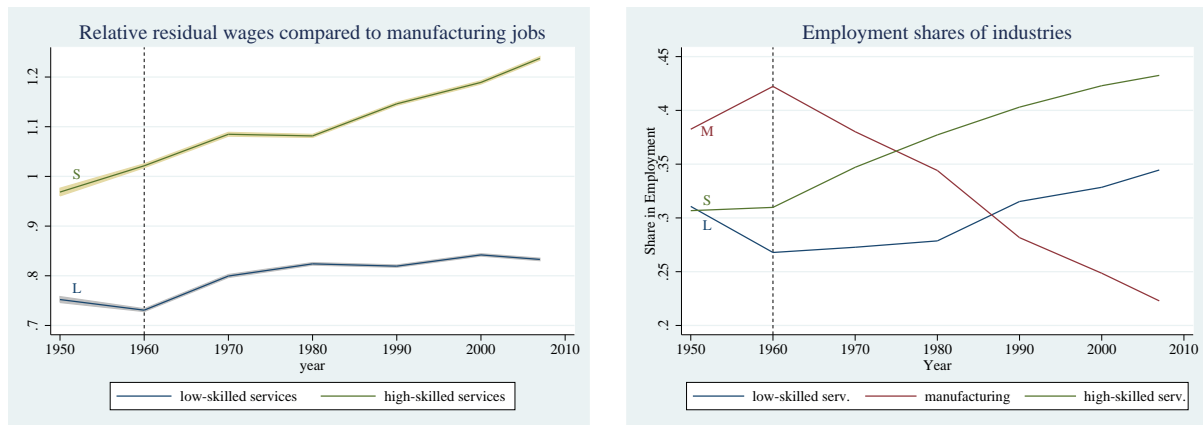


Figure 3: Polarization for broad industries

Notes: The data used is the same as in Figure 1. Each worker is classified into one of three sectors based on their industry code (for details of the industry classification see text and the appendix). The left 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. The right panel shows employment shares, calculated in terms of hours worked. The dashed vertical line represents 1960, from when on manufacturing employment has been contracting.

Figure 3 documents the patterns of polarization both in terms of employment shares and wages for the above defined sectors between 1950/1960 and 2007. The right panel shows the path of employment shares: high-skilled services increase continuously, low-skilled services increase and manufacturing decreases from 1960 onwards.¹⁶ In terms of wages, we plot the sector premium in high-skilled and low-skilled services compared to manufacturing, 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.¹⁷ We plot these rather than the relative average wages, 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

¹⁶Between 1950 and 1960 manufacturing employment increased, and low-skilled services dropped.

¹⁷See Table 6 in the appendix for details of the regressions. The patterns of relative average wages for industries are very similar, see Figure 13 in the appendix.

the differential path of these across sectors.¹⁸ As the graph shows, low-skilled service workers earn less, whereas high-skilled service workers earn more than manufacturing workers. Since the 1960s both low- and high-skilled service workers have been gaining in terms of wages compared to manufacturing workers. To summarize, from 1960 onwards 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 shifts

To quantify the contribution of sectoral employment shifts to each occupation's employment share path, we conduct a standard shift-share decomposition.¹⁹ 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 $\lambda_{oit} = L_{oit}/L_{it}$ denotes the share of occupation o , industry i employment within industry i employment at time t , and $E_{it} = L_{it}/L_t$ denotes the share of industry i employment within total employment at time t . 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

¹⁸One might be concerned that the employment share changes are driven by changes in the age, gender, race composition of the labor force. To assess this, we generate counterfactual industry employment shares by fixing the industry employment share of each age-gender-race cell at its 1960 level, and allowing the employment shares of the cells to change. This exercise confirms that to a large extent the employment share changes are not driven by the compositional changes of the labor force. See Figure 14 in the appendix.

¹⁹An alternative way is to calculate how much occupational employment shares would have changed, if industry employment shares would have remained at their 1960 level. See Figure 15 in the appendix.

sector forces.²⁰

Table 1: Decomposition of changes in occupational employment shares

	Employment shares			
	3 x 3		10 x 11	
	1950-2007	1960-2007	1950-2007	1960-2007
Manual				
Total Δ	2.98	5.68	2.98	5.68
Between Δ	2.30	3.07	3.13	4.38
Within Δ	0.67	2.61	-0.15	1.30
Routine				
Total Δ	-19.79	-19.14	-19.79	-19.14
Between Δ	-5.66	-6.32	-9.73	-10.01
Within Δ	-14.13	-12.82	-10.06	-9.13
Abstract				
Total Δ	16.81	13.46	16.81	13.46
Between Δ	3.35	3.24	6.60	5.63
Within Δ	13.46	10.21	10.21	7.83

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 over the period 1950-2007 and over 1960-2007. The first two columns use 3 occupations and 3 sectors, columns three and four 10 occupations and 11 industries.

Table 1 shows the occupational employment share changes and their decomposition between 1950 and 2007, and alternatively between 1960 and 2007, into a between-industry and a within-industry component. We show these changes for three broad occupational categories commonly used in the routinization literature. This Table shows that there has been polarization: employment has been shifting from routine to both manual and abstract jobs (also documented in Figure 12 in the appendix). In terms of average occupational employment share changes over this period, between 35 per cent and 53 per cent of the changes are driven by between-industry shifts, depending on whether we use a finer or a coarser categorization of occupations and industries. Between-industry shifts matter the most for manual occupations and the least for abstract occupations, where they still account for at least one fifth. The first two columns use the three occupation categories (manual, routine, abstract) and the three sectors

²⁰The 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 , $E_i = (E_{it} + E_{i0})/2$.

(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 finer categories, 10 occupations and 11 industry groups, shown in the last two columns.²¹

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 shift-share decomposition conducted here is very similar to that in Acemoglu and Autor (2011) and Goos et al. (2014), however, it is done for our classification of industries and occupations, our sample period and data. Both Acemoglu and Autor (2011) and Goos et al. (2014) found a similar magnitude as we do for the role of between-industry shifts.²² However, 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. While this is a valid concern, routinization, linked to ICT, is not likely to be driving the faster productivity growth observed between the 1950s and 1980s. We also conduct an alternative shift-share decomposition where we use sectoral value added shares instead of sectoral employment shares. Even though the importance of the within industry component seems to be smaller than in the standard shift-share decomposition, it is nonetheless a quite significant share of the overall change.²³ Our reading of these results is that in order to understand the occupational employment share changes it

²¹The 10 occupations are the same as in Figure 2, while the 11 industries are: 1 personal services; entertainment and low-skilled business and service repairs, 2 low-skilled transport, 3 retail trade, 4 wholesale trade, 5 extractive industries, 6 construction, 7 manufacturing, 8 professional and related services and high-skilled business services, 9 finance, insurance, and real estate, 10 high-skilled transport and public utilities (incl. communications), 11 public administration.

²²Acemoglu and Autor (2011) use US Census data between 1960 and 2000, and the ACS 2008 data to decompose occupational employment share changes. They aggregate their findings to four occupation categories (abstract, routine cognitive, routine non-cognitive, manual) and conduct the shift-share decomposition for men and women separately. Their focus is the declining importance of between-industry shifts from 1960-1980 to 1980-2007. We find some support for this in our decade-by-decade analysis shown in Table 8 in the appendix. The relatively smaller contribution of between-industry shifts in later periods might be due to routinization kicking in after the 1980s, thus providing an extra force for within-industry reallocation of labor. Nonetheless, we find that even in the 1980-2007 period, between-industry shifts explain a significant part of occupational employment share changes. Goos et al. (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.

²³See appendix for details.

is important to consider the forces that drive the structural shift of employment away from manufacturing and towards both types of services.

We also decompose relative occupational wage changes into an occupation driven component and an industry driven component. We define the relative average wage of occupation o as the ratio of the occupation's average wage relative to the average wage in routine occupations:²⁴

$$rw_{ot} \equiv \sum_i \frac{L_{iot}}{L_{ot}} \frac{w_{it}}{w_{rt}} \frac{w_{iot}}{w_{it}} = \sum_i \chi_{iot} rw_{it} p_{iot},$$

where $\chi_{iot} = L_{iot}/L_{ot}$ is employment in industry i and occupation o in period t relative to employment in occupation o in period t , $rw_{it} = w_{it}/w_{rt}$ is the ratio of the average wage in industry i in period t relative to the average wage of routine occupations in period t , and $p_{iot} = w_{iot}/w_{it}$ is the premium of occupation o in industry i in period t . For the decomposition it is also useful to define $rw_{iot} = w_{iot}/w_{rt}$, as the ratio of the average wage in industry i , occupation o in period t relative to the average wage of routine occupations in period t . We decompose the change in the relative wage of occupation o between period 0 and period t as follows:

$$\Delta rw_{ot} = rw_{ot} - rw_{o0} = \underbrace{\sum_i rw_{io} \Delta \chi_{iot} + \sum_i \chi_{io} p_{io} \Delta rw_{it}}_{\text{industry effect}} + \underbrace{\sum_i \chi_{io} rw_{io} \Delta p_{iot}}_{\text{occupation effect}},$$

where Δ denotes the change between period 0 and t , and the variables without a time subscript denote the average of the variable between period 0 and period t . The *industry effect* is itself composed of two parts: the first part captures that an occupation's average relative wage could change as workers within the occupation move to an industry where the occupation's relative average wage might be higher or lower. The

²⁴It is clear that the wage has to be a relative wage, otherwise the decomposition picks up the general upward trend in wages, and assigns it to the component which includes a wage change. If we were to normalize with the economy-wide average wage, that would capture some of the compositional effects of structural and occupational shifts. As we are mainly interested in the relative gain of manual and abstract wages compared to routine wages, we normalize with the routine wage. This comes at the cost that we cannot decompose the change in relative routine wages, as by construction, these do not change. Normalizing by the average wage of an industry would by construction restrict one of the industry components in the decomposition to be zero, as the relative average wage of that industry would not change.

second part captures that each industry’s relative average wage path influences the overall relative average wage of the occupation. The occupation effect is driven by the change in the occupational premium within each industry.

Table 2: Decomposition of changes in relative occupational wages

	Relative wages			
	3 x 3		10 x 11	
	1950-2007	1960-2007	1950-2007	1960-2007
Manual/Routine				
Total Δ	0.289	0.310	0.289	0.310
Industry Δ	0.181	0.148	0.225	0.218
Occupation Δ	0.107	0.162	0.064	0.093
Abstract/Routine				
Total Δ	0.327	0.240	0.327	0.240
Industry Δ	0.310	0.254	0.376	0.317
Occupation Δ	0.016	-0.013	-0.050	-0.077

Notes: Same data as in Figure 1. For each occupational category, the first row presents the total change, the second the industry component, and the third the occupation component over the period 1950-2007 and over 1960-2007. The first two columns use 3 occupations and 3 sectors, columns three and four 10 occupations and 11 industries.

Table 2 shows the change in manual and abstract wages relative to routine wages, and their decomposition into an industry and an occupation component, between 1950 and 2007 and 1960 and 2007. This Table confirms that both manual and abstract occupations have been gaining in terms of wages relative to routine jobs. The decomposition shows that between half and two thirds of the gain in relative manual wages is driven by industry effects, by the reallocation of manual labor to higher-paying industries, or by faster wage growth in industries where more manual workers are employed. For the increase in relative abstract wages the results are even more striking: all of the gain is driven by industry effects.

To summarize, we show two new facts and an additional one about the occupational and sectoral employment shares and relative wages. We documented that polarization defined over occupational categories 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 and relative wage changes in occupations is driven by the

(employment) shifts across industries.

In the rest of the paper we present a simple model of sorting and structural change to jointly explain the sectoral shifts in employment and the changes in average sectoral wages. We then calibrate the model to quantitatively assess how much of the polarization of sectoral employment and wages it can explain over the last fifty years, when feeding in sectoral labor productivity from the data.

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 labor and expenditure shares, but also about the relative average wages across sectors over time.

We assume that the economy is populated by heterogeneous agents, who all make individually optimal decisions about their sector of work. Every individual chooses their sector of work to maximize wages, 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 maximizes its utility subject to its budget constraint.²⁵ 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 wages, the stand-in household collects all wages and

²⁵We make the assumption of a stand-in household purely for expositional purposes. Given that the preferences we use are homothetic, the resulting sectoral demands are equal to the aggregation of individual demands.

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 1960s.

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:

$$Y_s = A_s N_s, \quad (1)$$

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

$$\omega_s = \frac{\partial p_s Y_s}{\partial N_s} = p_s A_s. \quad (2)$$

The technology to produce manufacturing goods is:

$$Y_m = A_m N_m, \quad (3)$$

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:

$$\omega_m = \frac{\partial p_m Y_m}{\partial N_m} = p_m A_m. \quad (4)$$

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

We assume that each worker is equally talented in providing low-skilled services, i.e. efficiency units of labor do not matter here. The total amount of low-skilled services provided is:

$$Y_l = A_l L_l, \quad (5)$$

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

$$\omega_l = \frac{\partial p_l Y_l}{\partial L_l} = p_l A_l. \quad (6)$$

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

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 wages 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 household 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 wages (as well as in all other members' wages), it is optimal for each worker to choose the sector which provides him with the highest wages.

Individuals are heterogeneous in their endowment of efficiency units of labor, $\mathbf{a} \in \mathbf{R}_+^2$, which is drawn from a time invariant distribution $f(\mathbf{a})$. This is an innate ability distribution, and as such is prior to any form of human capital that a worker might accumulate, for example by acquiring education. Even though we do not model this

explicitly, one could think of our model as having a reduced-form educational choice in the following sense. If a worker given their ability selects a sector, say S , to work in, they will on the way get education, or other qualifications, as necessary to enter that sector.²⁶

The endowment, \mathbf{a} , determines each individual's productivity in sector M and S , while each individual is equally productive in sector L . We assume that each dimension of ability corresponds to one sector, such 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.²⁷ Therefore the wage of an individual with $\mathbf{a} = (a_m, a_s)$ efficiency units of labor in sector L is ω_l , in sector M is $a_m\omega_m$, while if working in sector S it is $a_s\omega_s$.

Given wage rates $\omega_l, \omega_m, \omega_s$ – per unit of labor in sector L and per efficiency unit in M and S – the optimal decision of any agent can be characterized as follows.

Result 1. *Given unit wage rates ω_l, ω_m and ω_s , the optimal sector choice of individuals can be characterized by two cutoff values:*

$$\widehat{a}_m \equiv \frac{\omega_l}{\omega_m}, \quad (7)$$

$$\widehat{a}_s \equiv \frac{\omega_l}{\omega_s}. \quad (8)$$

It is optimal for an individual with (a_m, a_s) efficiency units of labor to work in sector L if and only if

$$a_m \leq \widehat{a}_m \quad \text{and} \quad a_s \leq \widehat{a}_s. \quad (9)$$

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

$$a_m \geq \widehat{a}_m \quad \text{and} \quad a_s \leq \frac{\widehat{a}_s}{\widehat{a}_m} a_m. \quad (10)$$

²⁶Given that high-skilled services is the most intensive sector in college educated workers, the increase in college education that the US has witnessed over the last decades is in line with our model, as it generates an increase in the demand for high-skilled services, providing a link from structural change to educational attainment (see for example also Buera, Kaboski, and Rogerson (2015)).

²⁷This assumption is without loss of generality. 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 $\alpha_m \neq \alpha_s$.

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

$$a_s \geq \hat{a}_s \quad \text{and} \quad a_m \leq \frac{\hat{a}_m}{\hat{a}_s} a_s \quad (11)$$

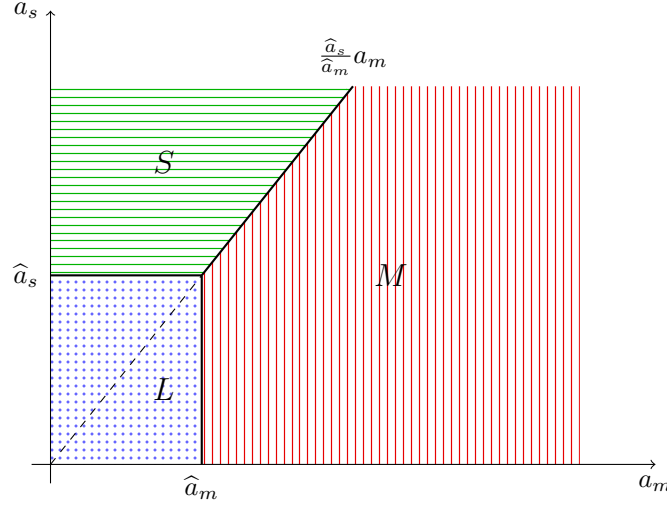


Figure 4: Optimal sector of work

The figure depicts the optimal sector of work choices as a function of $\hat{a}_m = \omega_l/\omega_m$ and $\hat{a}_s = \omega_l/\omega_s$. The blue dotted area shows the efficiency unit pairs (a_m, a_s) where L is the optimal sector, the red vertically striped area shows where M is optimal, and the green horizontally striped area shows where S is optimal.

Figure 4 shows this endogenous sorting behavior. Individuals who have low efficiency units in both manufacturing and high-skilled services sort into low-skilled services (the blue dotted area). Individuals with high enough manufacturing efficiency and relative to this a low high-skilled service efficiency sort into manufacturing jobs (the red vertically striped area). 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 horizontally striped area).

It is worth to consider the optimal sorting patterns as a function of relative unit wages. A ceteris paribus fall in the sector M unit wage, ω_m , makes working in sector M less attractive both compared to working in sector L and S . In the graph this change would be represented by an outward shift in $\hat{a}_m = \omega_l/\omega_m$, and a flattening of $a_m \hat{a}_s / \hat{a}_m = a_m \omega_m / \omega_s$; sector M is squeezed from both sides. A ceteris paribus fall in the sector S unit wage, ω_s , has similar effects. In the graph $\hat{a}_s = \omega_l/\omega_s$ would shift up, as sector L becomes more attractive, and $a_m \hat{a}_s / \hat{a}_m$ would become steeper as sec-

tor M becomes more attractive; sector S loses from both sides. A fall in sector L unit wage, ω_l , holding everything else constant, makes sector L less attractive. This would be captured in the graph by an inward shift of $\hat{a}_m = \omega_l/\omega_m$ and a downward shift of $\hat{a}_s = \omega_l/\omega_s$, leaving $a_m\hat{a}_s/\hat{a}_m$ unchanged; both sector M and S become more attractive compared to sector L .

The optimal sector of work choices of individuals determine the effective labor supplies in the three markets:

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

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

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

Note that in sector M and S these are the effective labor supplies, the raw labor supplies – or employment shares²⁸ – in these sectors are:

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

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

Note that $\omega_l, \omega_m, \omega_s$ are the sectoral unit wages, and we are also interested in the sectoral average wages. These simply are the total earnings in a sector divided by the mass of people working in the sector:

$$\bar{w}_l = \frac{\int_0^{\hat{a}_m} \int_0^{\hat{a}_s} \omega_l f(a_m, a_s) da_s da_m}{L_l} = \omega_l, \quad (17)$$

$$\bar{w}_m = \frac{\int_{\hat{a}_m}^{\infty} \int_0^{\frac{\hat{a}_s}{\hat{a}_m} a_m} \omega_m a_m f(a_m, a_s) da_s da_m}{L_m} = \omega_m \frac{N_m}{L_m}, \quad (18)$$

$$\bar{w}_s = \frac{\int_{\hat{a}_s}^{\infty} \int_0^{\frac{\hat{a}_m}{\hat{a}_s} a_s} \omega_s a_s f(a_m, a_s) da_m da_s}{L_s} = \omega_s \frac{N_s}{L_s}. \quad (19)$$

²⁸The empirical counterpart of these raw labor supplies, L_l, L_m, L_s are the sectoral employment shares, as $L_l + L_m + L_s = 1$.

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:

$$\begin{aligned} \max_{C_l, C_m, C_s} u \left(\left[\theta_l C_l^{\frac{\varepsilon-1}{\varepsilon}} + \theta_m C_m^{\frac{\varepsilon-1}{\varepsilon}} + \theta_s C_s^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{\varepsilon}{\varepsilon-1}} \right) \\ \text{s.t.} \quad p_l C_l + p_m C_m + p_s C_s \leq \omega_l L_l + \omega_m N_m + \omega_s N_s \end{aligned}$$

where u is any monotone increasing function, $\omega_l L_l + \omega_m N_m + \omega_s N_s$ are the total wages of household members, 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:

$$\frac{C_l}{C_m} = \left(\frac{p_l \theta_m}{p_m \theta_l} \right)^{-\varepsilon}, \quad (20)$$

$$\frac{C_s}{C_m} = \left(\frac{p_s \theta_m}{p_m \theta_s} \right)^{-\varepsilon}. \quad (21)$$

3.3 Competitive equilibrium and structural change

A competitive equilibrium is given by cutoff sector-of-work efficiencies $\{\hat{a}_m, \hat{a}_s\}$, unit wage rates $\{\omega_l, \omega_m, \omega_s\}$, 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 $i = L, M, S$), where the supply is given by (1), (3) and (5), and the market clearing unit wage rates, (2), (4) and (6), in the household's optimality conditions, (20) and (21), we obtain the following:

$$\begin{aligned} \frac{A_l}{A_m} \frac{L_l}{N_m} &= \left(\underbrace{\frac{\omega_l}{\omega_m} \frac{A_m}{A_l}}_{=p_l/p_m} \frac{\theta_m}{\theta_l} \right)^{-\varepsilon}, \\ \frac{A_s}{A_m} \frac{N_s}{N_m} &= \left(\underbrace{\frac{\omega_s}{\omega_m} \frac{A_m}{A_s}}_{=p_s/p_m} \frac{\theta_m}{\theta_s} \right)^{-\varepsilon}. \end{aligned}$$

The left hand side is the relative supply, while the right hand side is the relative demand for low- and respectively high-skilled services compared to manufacturing. A change in the relative productivity affects both the relative supply and the relative demand.

An increase in relative manufacturing productivity compared to low-skilled service productivity (A_m/A_l) has two direct effects: (i) it reduces the relative supply of low-skilled services, (Y_l/Y_m), and (ii) through an increase in the relative price of low-skilled services (p_l/p_m), it lowers the relative demand for low-skilled services. If low-skilled services and manufacturing goods are complements, $\varepsilon < 1$, the effect through relative prices is the weaker one, and relative supply falls by more than relative demand. To restore equilibrium, the relative supply of low-skilled services has to increase and/or its relative demand has to fall compared to manufacturing.

In order for the relative supply to increase, the efficiency units of labor hired in low-skilled services have to increase relative to manufacturing, which requires a rise in the relative unit wage, ω_l/ω_m . At the same time, a rise in the relative unit wage also increases the relative price of low-skilled services, thus lowering the relative demand. Therefore the equilibrium requires a rise in the relative low-skilled service unit wage compared to the manufacturing unit wage. Similarly, an increase in A_m/A_s , through its affect on relative supply and relative demand, requires a rise in ω_s/ω_m .

Using the definition of the optimal sector-of-work cutoffs (7) and (8), we can express the relative unit wages as $\omega_l/\omega_m = \hat{a}_m$ and $\omega_s/\omega_m = \hat{a}_m/\hat{a}_s$. Substituting this into the two above equations, and using the optimal sorting of individuals, (12), (13) and (14), we obtain the following expressions, which allow us to formally analyze the comparative static properties of the equilibrium:

$$\frac{L_l(\hat{a}_m, \hat{a}_s)}{N_m(\hat{a}_m, \hat{a}_s)} \hat{a}_m^\varepsilon = \left(\frac{A_m}{A_l} \right)^{1-\varepsilon} \left(\frac{\theta_m}{\theta_l} \right)^{-\varepsilon}, \quad (22)$$

$$\frac{N_s(\hat{a}_m, \hat{a}_s)}{N_m(\hat{a}_m, \hat{a}_s)} \left(\frac{\hat{a}_m}{\hat{a}_s} \right)^\varepsilon = \left(\frac{A_m}{A_s} \right)^{1-\varepsilon} \left(\frac{\theta_m}{\theta_s} \right)^{-\varepsilon}. \quad (23)$$

These two equations implicitly define the equilibrium sector-of-work cutoffs, \hat{a}_m and \hat{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 $\hat{a}_m = \omega_l/\omega_m$ and $\hat{a}_m/\hat{a}_s = \omega_s/\omega_m$ unambiguously increase, while $\hat{a}_s = \omega_l/\omega_s$ can rise or fall. This results in an unambiguous increase in labor in L , in efficiency labor in S , and a reduction in efficiency and raw labor in M .*

Proof. Total differentiation of (22) and (23). See appendix for details. \square

Proposition 1 confirms the results of Ngai and Pissarides (2007) in terms of efficiency labor, rather than raw labor or employment shares: when sectoral outputs are complements in consumption, efficiency labor needs to reallocate to the sectors which become relatively less productive. As manufacturing productivity grows the fastest, efficiency labor has to move out of manufacturing into both low- and high-skilled services. Since individuals optimally sort into the sector with the highest return for them, this implies that the optimal sector-of-work cutoffs have to adjust. Proposition 1 states and Figure 5 depicts what these adjustments entail. The adjustment to the new equilibrium requires sector M to be squeezed from both sides, $\hat{a}_m = \omega_l/\omega_m$ has to shift to the right (ω_l has to increase relative to ω_m), and $a_m\hat{a}_s/\hat{a}_m = a_m\omega_m/\omega_s$ has to become flatter (ω_s has to increase relative to ω_m). This is very intuitive: sector M has to shrink, while sector L and S have to expand, which requires sector M unit wages to fall both relative to sector S and sector L unit wages. It is worth to note that these results hold for any underlying distribution of efficiency units of labor, $f(a_m, a_s)$. However, since sector L and S productivity grow at the same rate, in general it is ambiguous whether the boundary between L and S shifts up or down ($\hat{a}_s = \omega_l/\omega_s$ increases or decreases).

To understand what these adjustments imply for employment shares (or for raw labor) Figure 5 is useful. Efficiency labor in M is the aggregate amount of a_m in the area bounded by \hat{a}_m and $a_m\hat{a}_s/\hat{a}_m$, while raw labor is just the mass of individuals in the same area. Both the outward shift of \hat{a}_m and the flattening of \hat{a}_s/\hat{a}_m have a negative impact on both raw and efficiency labor in M . Similarly efficiency labor in S is the aggregate amount of a_s in the area bounded by \hat{a}_s and $a_m\hat{a}_s/\hat{a}_m$, while raw labor is the mass in the same area. If \hat{a}_s falls, as shown in the left panel, then efficiency and raw

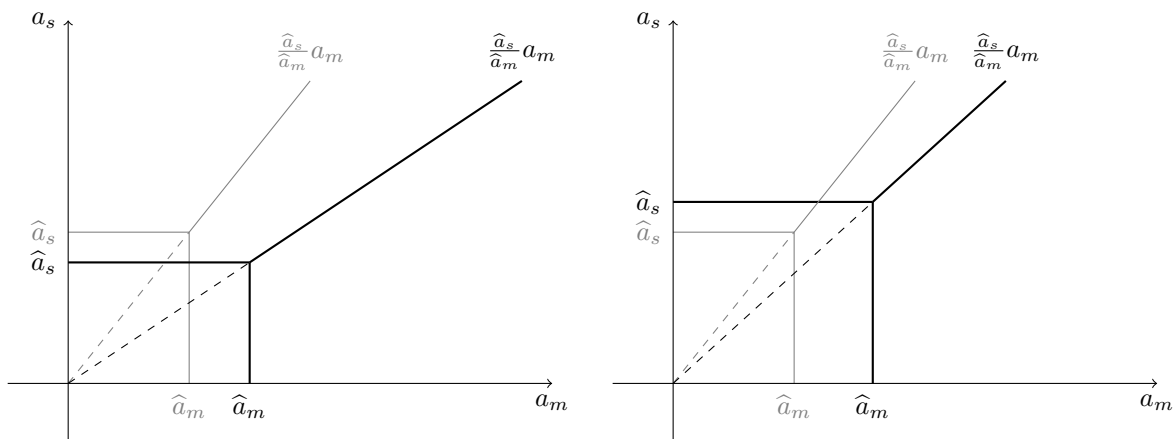


Figure 5: Change in the optimal sorting

The figure shows the possible shifts in \hat{a}_m and \hat{a}_s in response to an increase in relative manufacturing productivity when the sectoral outputs are complements in consumption. The original cutoffs are shown in gray, while the new ones are shown in black and bold. While \hat{a}_m unambiguously increases, and \hat{a}_s/\hat{a}_m unambiguously decreases, the direction of change in \hat{a}_s is ambiguous.

labor in S increases. If \hat{a}_s increases, shown in the right panel, raw labor in S might fall, while efficiency labor in S has to increase. The fact that efficiency labor in S has to increase implies that the amount of efficiency labor in S lost due to the increase in \hat{a}_s is smaller than the amount gained due to the decrease in \hat{a}_s/\hat{a}_m . However, since the average a_s in the area lost is strictly lower than in the area gained, we cannot determine, in general, whether in S raw labor increases or falls. Since in low-skilled services efficiency and raw labor units are the same, the employment share of low-skilled services increases.

In terms of relative average wages, in general it is not possible to sign the changes predicted by the model. The reason is self-selection; the workers leaving manufacturing are the ones that have a relatively low efficiency. As a consequence, the average efficiency in manufacturing increases when its employment share decreases. This tends to increase the average wage in manufacturing compared to the other sectors, offsetting to some extent the direct effect of the falling relative manufacturing unit wage. Without further assumptions, it is conceivable that the indirect effect through the average efficiency might overturn the direct effect of changing unit wages. To see this, consider the average low-skilled service wage, (17), relative to the average manufacturing wage, (18):

$$\frac{\bar{w}_l}{\bar{w}_m} = \frac{\omega_l}{\omega_m \frac{N_m}{L_m}} = \frac{\hat{a}_m}{\bar{a}_m},$$

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

$$\frac{d\frac{\bar{w}_l}{\bar{w}_m}}{\frac{\bar{w}_l}{\bar{w}_m}} = \frac{d\hat{a}_m}{\hat{a}_m} - \frac{d\bar{a}_m}{\bar{a}_m},$$

i.e. the difference between the percentage change in the sector-of-work cutoff between L and M and the percentage change in the average efficiency in sector M . From Proposition 1 we know that $\frac{d\hat{a}_m}{\hat{a}_m}$ is positive, the question is whether the average efficiency in M can increase more in percentage terms than this.

Similarly the average high-skilled service wage, (19), relative to the average manufacturing wage can be expressed as:

$$\frac{\bar{w}_s}{\bar{w}_m} = \frac{\omega_s \frac{N_s}{L_s}}{\omega_m \frac{N_m}{L_m}} = \frac{\hat{a}_m \bar{a}_s}{\hat{a}_s \bar{a}_m},$$

where $\bar{a}_s \equiv N_s/L_s$ is the average efficiency of high-skilled service workers in S . This implies that the percentage change in the relative wage is:

$$\frac{d\frac{\bar{w}_s}{\bar{w}_m}}{\frac{\bar{w}_s}{\bar{w}_m}} = \frac{d\hat{a}_m}{\hat{a}_m} - \frac{d\hat{a}_s}{\hat{a}_s} + \frac{d\bar{a}_s}{\bar{a}_s} - \frac{d\bar{a}_m}{\bar{a}_m}.$$

The direct effect of the change in the cutoffs ($d\hat{a}_m/\hat{a}_m - d\hat{a}_s/\hat{a}_s$) is again positive by Proposition 1, while in general its indirect effect through the average sectoral abilities goes in the opposite direction. Again the overall change cannot be signed independent of the underlying distribution of efficiency units and of the initial level of labor in the different sectors. Even though in general the direction of wage changes might depend on the form of the underlying distribution of abilities, for the most commonly used classes of distributions in all our simulations the relative average wages (for both L and S to M) moved in the same direction as the cutoffs.

Since the structural change literature focuses on employment shares and value added shares, we also investigate our model's implications for the relative value added. 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:*

$$d \frac{p_s Y_s}{p_m Y_m} > 0 \quad \text{and} \quad d \frac{p_l Y_l}{p_m Y_m} > 0.$$

These results can be understood by considering the following. In this model, sectoral value added is equal to the sectoral wage bill: $p_i Y_i = p_i A_i N_i = \omega_i N_i$. Proposition 1 tells us that $\omega_l/\omega_m = \hat{a}_m$ increases, that L_l increases, and N_m falls. Both relative unit wages and effective labor changes increase the value added output of sector L relative to sector M . Similarly, $\omega_s/\omega_m = \hat{a}_m/\hat{a}_s$ also increases according to Proposition 1, while efficiency labor in S increases and in M it falls.

The sectoral value added can be further expressed as $p_i Y_i = \omega_i N_i = \bar{w}_i L_i$, since the sectoral wage bill can be expressed as either sectoral unit wage times sectoral efficiency labor, or as sectoral average wage times sectoral raw labor. Using this latter expression we can show that

$$\frac{p_i Y_i}{p_j Y_j} = \frac{\bar{w}_i L_i}{\bar{w}_j L_j}.$$

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

3.4 Calibration

The four key moments are the relative average sectoral wages, $\frac{\bar{w}_l}{\bar{w}_m}$ and $\frac{\bar{w}_s}{\bar{w}_m}$, and the sectoral employment shares, L_l , L_m and L_s , which sum to one. Data for the average sectoral wages and the sectoral employment shares come from the 1960 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 $\theta_l, \theta_m, \theta_s, \varepsilon$, the distribution of labor efficiencies, $f(a_m, a_s)$, and the initial sectoral labor productivities, $A_l(0), A_m(0), A_s(0)$.

We present a very simple model of structural change, and we are only interested in its predictions for sectoral employment shares and relative average sectoral wages, but not in its predictions for any higher moment of the wage distribution. We therefore choose a class of distributions for the labor efficiency endowments which requires only a minimal choice of parameters: the (bivariate) uniform distribution. While we established in Proposition 1 of section 3.3 the main qualitative predictions of the model, it is likely that the quantitative effects depend on the class of underlying distributions of sectoral efficiencies. For this reason, in section 3.6 we show our model's predictions for various alternative distributions, the parametrization of which require additional assumptions.

For the baseline calibration we assume that labor efficiencies are distributed uniformly. Without loss of generality we normalize the mean of a_m and a_s to be unity.²⁹ Given these assumptions, two parameters of the distribution are left to be calibrated, the minimum (and thus the maximum, given that the mean is one) of both labor efficiencies, denoted by \underline{a}_m and \underline{a}_s . We calibrate these two parameters of the distribution of efficiency units such that the observed employment shares and the relative average wages in 1960 are consistent with each other, in the sense that if the model matches employment, it also matches relative average wages. The procedure that we implement is the following. For any given distribution, we can compute what the observed

²⁹See appendix for details that this is indeed a normalization.

employment shares imply for average relative wages. In particular, for any \underline{a}_m and \underline{a}_s , we can compute the cutoffs \hat{a}_m and \hat{a}_s that are implied by the observed employment shares, i.e. the $L_l(\hat{a}_m, \hat{a}_s)$, $L_m(\hat{a}_m, \hat{a}_s)$, $L_s = 1 - L_l(\hat{a}_m, \hat{a}_s) - L_m(\hat{a}_m, \hat{a}_s)$ seen in the 1960 data. These values for $\underline{a}_m, \underline{a}_s$ and \hat{a}_m, \hat{a}_s imply for relative average wages (combining (17),(18) and (19)):

$$\begin{aligned}\frac{\bar{w}_l}{\bar{w}_m} &= \frac{\omega_l}{\omega_m \frac{N_m(\hat{a}_m, \hat{a}_s)}{L_m(\hat{a}_m, \hat{a}_s)}} = \hat{a}_m \frac{L_m(\hat{a}_m, \hat{a}_s)}{N_m(\hat{a}_m, \hat{a}_s)}, \\ \frac{\bar{w}_s}{\bar{w}_m} &= \frac{\omega_s \frac{N_s}{L_s}}{\omega_m \frac{N_m}{L_m}} = \frac{\hat{a}_m}{\hat{a}_s} \frac{L_m(\hat{a}_m, \hat{a}_s) N_m(\hat{a}_m, \hat{a}_s)}{L_s(\hat{a}_m, \hat{a}_s) N_s(\hat{a}_m, \hat{a}_s)}.\end{aligned}$$

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. Choosing the parameters of the efficiency units distribution in this way guarantees that if we match the employment shares, then the relative average wages in 1960 are matched as well, and vice versa. However, we still have to calibrate other parameters to ensure that in equilibrium we are matching these moments in the first place. For this, we can either target the relative average wages, the employment shares, or the sector-of-work cutoffs implied by these.

Left to calibrate are the parameters of the utility function and the initial labor productivities. Previous literature has found a very low elasticity of substitution between goods and services when output is measured in consumption value added terms. Ngai and Pissarides (2008) find that plausible estimates are in the range (0, 0.3), while Herrendorf et al. (2013) find a value of $\varepsilon = 0.002$, which we use in our baseline calibration. Of the remaining six parameters, $\theta_l, \theta_m, \theta_s$ and $A_l(0), A_m(0), A_s(0)$ only two ratios matter for the equilibrium of this economy, as can be seen from (22) and (23). It is only the ratio of the θ s to the power of $-\varepsilon$ multiplied by the ratios of initial labor productivities to the power of $1 - \varepsilon$ that matter. Thus we calibrate $\tau_l \equiv \left(\frac{A_m(0)}{A_l(0)}\right)^{1-\varepsilon} \left(\frac{\theta_m}{\theta_l}\right)^{-\varepsilon}$ and $\tau_s \equiv \left(\frac{A_m(0)}{A_s(0)}\right)^{1-\varepsilon} \left(\frac{\theta_m}{\theta_s}\right)^{-\varepsilon}$ to match the sector-of-work cutoffs in 1960. This guarantees that both the relative average sectoral wages and the employment shares in 1960 are matched. The calibrated parameters are summarized in Table 3.

Similarly to Ngai and Petrongolo (2014) we calculate labor productivity growth by dividing sectoral value added output data from Herrendorf et al. (2013) with sectoral

Table 3: Calibrated Parameters

	Description	Value
$[a_m, \tilde{a}_m]$	range of manufacturing efficiency	[0.40, 1.60]
$[a_s, \tilde{a}_s]$	range of high-skilled service efficiency	[0.02, 1.98]
ε	CES b/w L , M and S in consumption	0.002
τ_l	relative weight on M	0.49
τ_s	relative weight on S	0.91

The value of the elasticity of substitution, ε , is taken from the literature. Conditional on its value, the remaining parameters are chosen to match the sectoral employment shares and relative average wages in 1960. In the robustness checks, we recalibrate the τ 's for different values of ε .

labor data from the Bureau of Economic Analysis (BEA). We rely on Herrendorf et al. (2013) for the consumption value added data rather than taking these data directly from the BEA, since as Herrendorf et al. 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 neither the industry level data on consumption value added output in Herrendorf et al. (2013), nor the employment data from the BEA is detailed enough, 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.³⁰

Table 4: Annual average labor productivity growth

	Based on raw labor		Adjusted by average efficiency	
	Manufacturing	Services	Manufacturing	Services
1960-1970	1.0220	1.0130	1.0210	1.0137
1970-1980	1.0155	1.0078	1.0145	1.0085
1980-1990	1.0304	1.0060	1.0277	1.0064
1990-2000	1.0316	1.0143	1.0303	1.0149
2000-2007	1.0263	1.0143	1.0245	1.0146
1960-2007	1.0251	1.0109	1.0235	1.0114

Labor productivity growth rates in the first two columns are calculated by dividing sectoral value added data taken from Herrendorf et al. (2013) 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 labor efficiencies.

Table 4 contains the average annual labor productivity growth in manufacturing and in low- and high-skilled services jointly for each decade between 1960 and 2007,

³⁰See appendix for details.

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, we evaluate the quantitative performance of the model in two ways: (i) by plugging in the average annual growth rates for the entire period; and (ii) by plugging in the decennial growth rates. The first two columns contain these growth rates when using raw employment growth.

It is well known that if individuals self-select based on their endowments of efficiency units and one cannot observe these efficiency units, then the measurement of changes in average wages or in labor productivity will be biased.³¹ In our model, expanding sectors soak up, while contracting sectors shed relatively less efficient workers. This implies that the average efficiency in manufacturing increases, while the average efficiency in high-skilled services falls over time, which – if left uncorrected – leads to an overestimation of productivity growth in manufacturing relative to services.³² To understand the potential magnitude of this bias, we calculate the change in average sectoral labor efficiencies implied by the change in the cutoff efficiencies required to match sectoral employment shares between 1960 and 2007, for the calibrated distribution of underlying labor efficiencies. In Table 4 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 efficiency units distribution, the bias in the relative productivity differentials is very small.³³

3.5 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 prefer-

³¹Carneiro 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.

³²This is true only when manufacturing employment is shrinking, and services employment is expanding.

³³Young (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 labor efficiencies the bias is relatively small.

ence and the sectoral efficiency distribution parameters at the values calibrated using data only from 1960, and feed in various labor productivity growth measures as calculated in Table 4. Our ultimate interest is the endogenous path of employment shares and relative average wages.

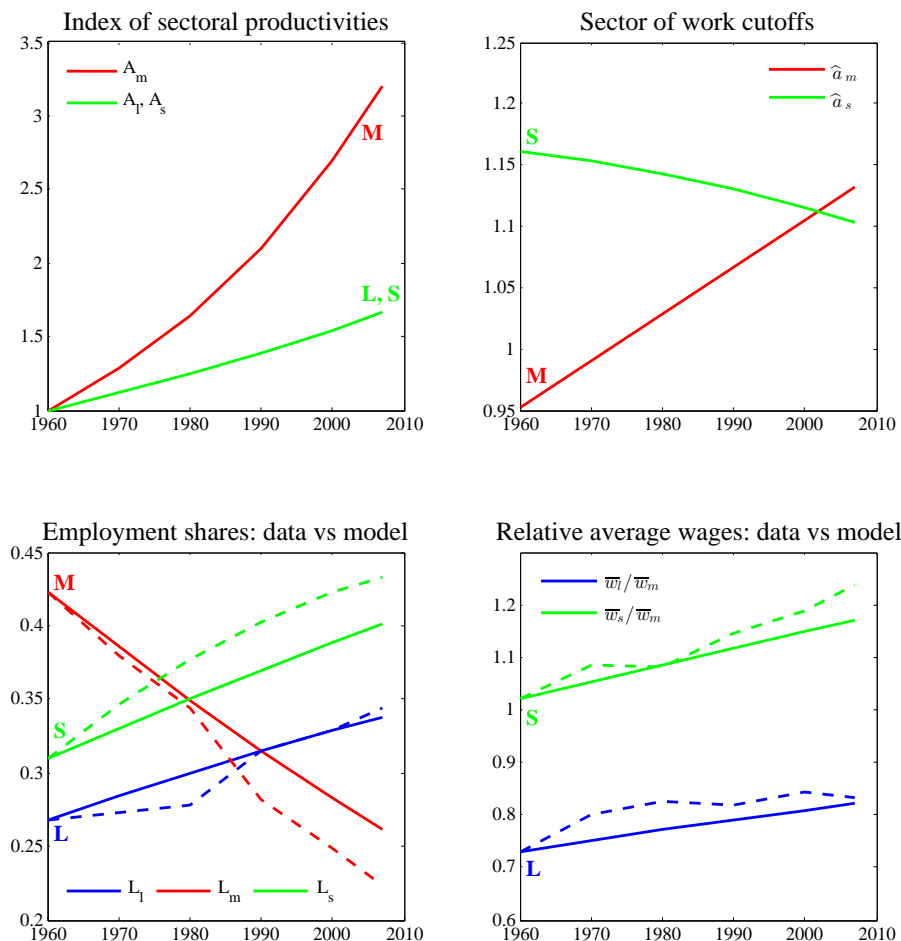


Figure 6: Transition of the benchmark model

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 6 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 1960-2007: 2.51% annual growth in manufacturing, and 1.09% growth in combined services (bottom left numbers in Table 4). The top left panel shows the path of M and of both L and S sector productivity. 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 labor in low- and high-skilled services puts an upward pressure on the unit wages in these sectors 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 \hat{a}_m and \hat{a}_s : the cutoff efficiency between L and M increases, while the cutoff between L and S decreases monotonically. This implies a continuous increase in S relative to L sector unit wages, which improve relative 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 1960 employment shares (bottom left panel) and the 1960 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 employment shares and relative average wages after 1960. Our baseline model predicts at least three quarters of the change in the employment share of each sector.³⁴ 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 labor efficiencies. As mentioned earlier, these two effects, in general, go in opposite directions, however the direct effect of the unit wages typically dominates the indirect effect that it has on average sectoral efficiencies. This is true in our baseline model as well, and our model overall predicts about 90 per cent of the growth in the relative low-skilled service sector wages, and 70 per cent of the growth in the relative average high-skilled service sector wages compared to manufacturing.³⁵

As we argued earlier, the differences in productivity growth between services and manufacturing are in general overestimated when not taking into account the effects

³⁴In the data the employment share of the high-skilled service sector increased by 12 percentage points, our model predicts a 9 percentage point increase. The employment share of manufacturing workers in the data fell by 20 percentage points, our model predicts a 16 percentage point contraction. Finally, the low-skilled service sector employment share increased by 8 percentage points in the data, whereas our model predicts a 7 percentage point increase.

³⁵In the data the relative average wage of low-skilled service workers compared to manufacturing workers increased by 14 per cent, while that of the high-skilled service workers increased by 21 per cent. In the simulation the average wage in low-skilled services compared to manufacturing increased by 12 per cent (19 per cent improvement in relative unit wages and 6 per cent decline in relative average efficiency), while the relative average high-skilled service sector wages increased by 14 per cent (25 per cent increase in relative unit wages and a more than 8 per cent drop in relative average efficiency).

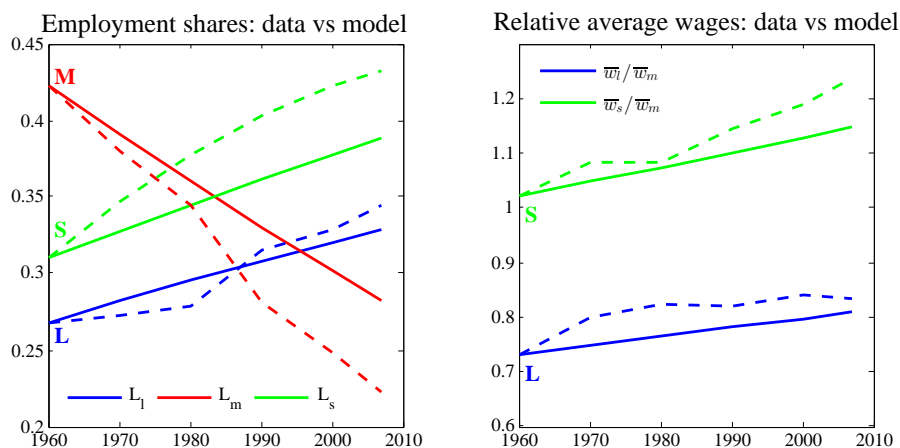


Figure 7: Transition of the model with selection-adjusted productivity growth rates

of the endogenous sorting of workers, which systematically changes the average efficiency in each sector. We therefore repeat our numerical exercise with the baseline calibration, but feeding in the average productivity growth rates adjusted for this bias. As can be seen in Table 4, when correcting for selection, we get a slightly lower 2.35% average annual productivity growth in manufacturing, and a slightly higher 1.14% annual rate in services. Figure 7 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 we correct for the changing selection of individuals, the equilibrium requires less labor to shift out of manufacturing. This brings the model's predictions a bit further away from the data, both in terms of employment shifts and relative wages.

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 1960 and 2007. However, Table 4 reveals that the growth rates have varied substantially over time. Figure 8 shows the simulated model contrasted with the data when feeding in the growth rates calculated for each period using raw employment.³⁶ 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

³⁶Figure 16 in the appendix shows the model's predictions for the decennial selection-adjusted growth rates.

are the same. Quantitatively, when feeding in the actual series of productivity growth (shown in the top left panel), the model predicts the actual time path of these variables even better. In the data manufacturing productivity growth accelerated relative to services in 1980, which might be driven by routinization having a stronger impact in manufacturing. 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 expansion and manufacturing employment contraction accelerated.

Since employment and wages are the focus of this paper, we targeted their 1960 level in our calibration. As we discussed in section 3.3, a model without capital intensity differences and with perfect capital mobility across sectors cannot match the level of sectoral relative average wages, employment and expenditure shares jointly.

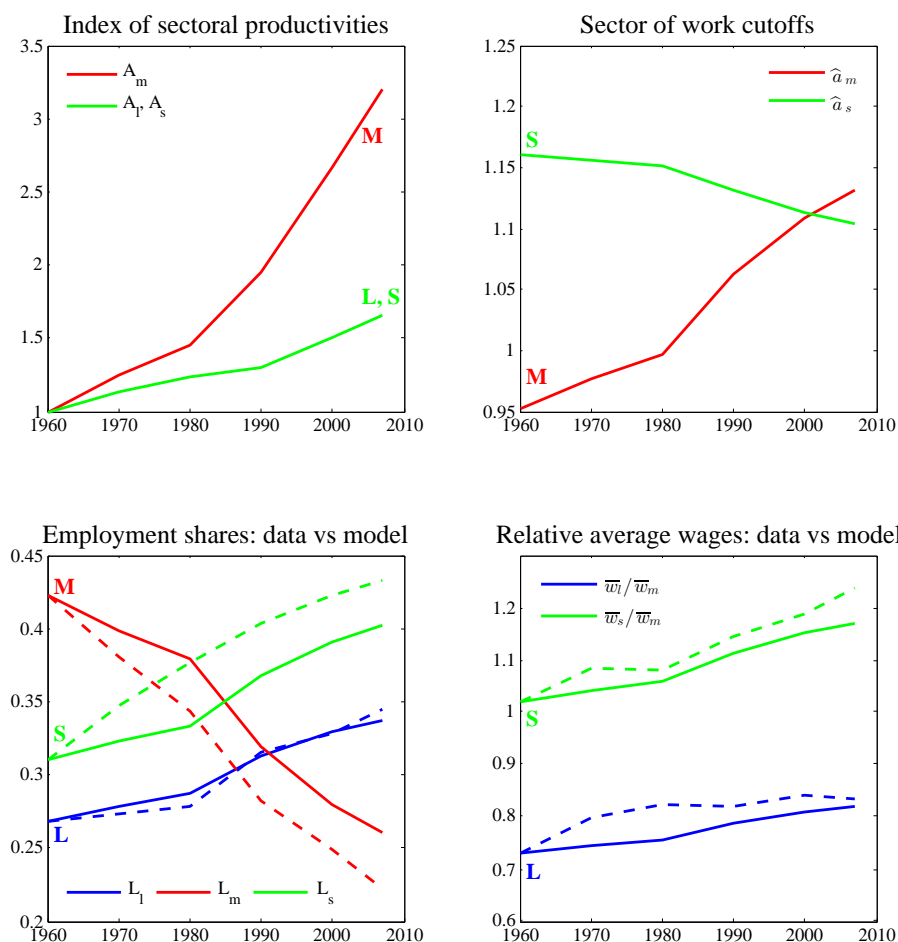


Figure 8: Transition of the model with decennial growth

Nonetheless, our model does quite a good job in predicting the path of the relative value added share in manufacturing compared to combined services between 1960 and 2007. Figure 9 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).³⁷ Even though the level of relative value added is not matched by the model, the overall decline is replicated quite well: in the data it declined by 53 per cent, while in the model it declined by 58 per cent.

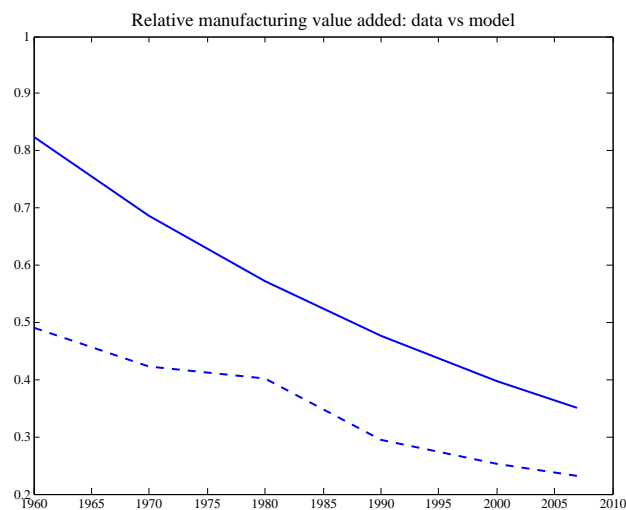


Figure 9: Transition of relative manufacturing value added

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).

3.6 Robustness checks

In our baseline calibration we assume that the distribution of sectoral efficiencies is uniform. As discussed earlier, while our main qualitative predictions do not depend on the distribution, the quantitative predictions are likely to be affected by it. In particular, assuming a less dispersed and more correlated distribution takes the model closer to one where labor is homogeneous (i.e. similar to Ngai and Pissarides (2007)). In the limiting case, where all individuals are endowed with the same amount of efficiency units of labor in all sectors, unit and average sectoral wages need to be equalized at all times, implying constant relative average wages. We recalibrate the model assuming

³⁷We 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.

a normal distribution (truncated at zero for both a_m and a_s) and a lognormal distribution for three levels of correlation: 0, 0.3 and 0.5. As Table 10 in the appendix shows, the model's quantitative predictions for employment share changes are virtually unchanged in all cases. While average wages relative to manufacturing increase in both the high- and the low-skilled service sector for all calibrations, the magnitude of these changes varies quite a bit. As we impose a higher correlation between a_m and a_s , matching employment shares and relative wages in 1960 requires a more compressed distribution. The higher the correlation and the lower the variance of a_m and a_s , the smaller is the predicted adjustment in the relative average sectoral wages. This is not surprising as the model gets closer to the case of homogeneous labor. In the cases considered in Table 10, even with the highest correlation and smallest variances, the model predicts one third of the increase the average low-skilled and one seventh of the increase in average high-skilled service sector wage relative to the average manufacturing wage.

In our baseline calibration we use $\varepsilon = 0.002$ for the elasticity of substitution between goods and services (measured in value-added terms) as estimated by Herrendorf et al. (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 1960 data.³⁸ Qualitatively the transition paths look exactly the same. A higher 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 employment shares and in relative average wages. Increasing the value of the elasticity of substitution takes the model's predictions further away from the time paths observed in the data. But as can be seen in Table 11 in the appendix, even with the least favorable calibration (high ε and labor productivity growth adjusted for selection), our model predicts 59 per cent of the increase in L and 44 per cent of the increase in S sector average wages relative to M . In terms of employment share changes, the model predicts

³⁸For $\varepsilon = 0.02$ these are $\tau_l = 0.49$ and $\tau_s = 0.91$; for $\varepsilon = 0.2$ we obtain $\tau_l = 0.48$ and $\tau_s = 0.88$.

at least half of the observed changes between 1960 and 2007.³⁹ Overall, the benchmark calibration with $\varepsilon = 0.002$ as estimated by Herrendorf et al. (2013) seems to do best in replicating the data.

4 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. This simple model does remarkably well in predicting the sectoral wage and employment patterns of the last 50 years.

³⁹It predicts a 5 percentage point increase in L , an 11 percentage point fall in M , and a 6 percentage point increase in S employment share compared to the 8, 20 and 12 percentage point changes in the data.

References

- Daron Acemoglu and David Autor. Chapter 12 - Skills, Tasks and Technologies: Implications for Employment and Earnings. In Orley Ashenfelter and David Card, editors, *Handbook of Labor Economics*, volume 4, Part B, pages 1043 – 1171. Elsevier, 2011.
- Daron Acemoglu and Veronica Guerrieri. Capital Deepening and Nonbalanced Economic Growth. *Journal of Political Economy*, 116(3):pp. 467–498, 2008.
- David H. Autor and David Dorn. The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–97, 2013.
- David H. Autor, Frank Levy, and Richard J. Murnane. The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):pp. 1279–1333, 2003.
- David H. Autor, Lawrence F. Katz, and Melissa S. Kearney. The Polarization of the U.S. Labor Market. *The American Economic Review*, 96(2):pp. 189–194, 2006.
- Alan S. Blinder. How Many US Jobs Might Be Offshorable? *World Economics*, 10(2): 41–78, 2009.
- Alan S. Blinder and Alan B. Krueger. Alternative Measures of Offshorability: A Survey Approach. *Journal of Labor Economics*, 31(S1):S97 – S128, 2013.
- Timo Boppart. Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences. *Econometrica*, forthcoming, 2014.
- Francisco J. Buera and Joseph P. Kaboski. The Rise of the Service Economy. *American Economic Review*, 102(6):2540–69, September 2012.
- Francisco J. Buera, Joseph P. Kaboski, and Richard Rogerson. Skill biased structural change. Working Paper 21165, National Bureau of Economic Research, May 2015.
- Pedro Carneiro and Sokbae Lee. Trends in Quality-Adjusted Skill Premia in the United States, 1960-2000. *American Economic Review*, 101(6):2309–49, 2011.

- Francesco Caselli and Wilbur John Coleman, II. The U.S. Structural Transformation and Regional Convergence: A Reinterpretation. *Journal of Political Economy*, 109(3): pp. 584–616, 2001.
- David Dorn. *Essays on Inequality, Spatial Interaction, and the Demand for Skills*. PhD thesis, University of St. Gallen, September 2009. Dissertation no. 3613.
- Christian Dustmann, Johannes Ludsteck, and Uta Schönberg. Revisiting the German Wage Structure. *The Quarterly Journal of Economics*, 124(2):843–881, 2009.
- Andy Feng and Georg Graetz. Rise of the Machines: The Effects of Labor-Saving Innovations on Jobs and Wages. Working paper, 2014.
- Maarten Goos and Alan Manning. Lousy and Lovely Jobs: The Rising Polarization of Work in Britain. *The Review of Economics and Statistics*, 89(1):118–133, February 2007.
- Maarten Goos, Alan Manning, and Anna Salomons. Job Polarization in Europe. *The American Economic Review Papers and Proceedings*, 99(2):pp. 58–63, May 2009.
- Maarten Goos, Alan Manning, and Anna Salomons. Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*, 104(8):2509–26, 2014.
- Gene M. Grossman and Esteban Rossi-Hansberg. Trading Tasks: A Simple Theory of Offshoring. *American Economic Review*, 98(5):1978–97, 2008.
- Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi. Two Perspectives on Preferences and Structural Transformation. *American Economic Review*, 103(7):2752–89, 2013.
- Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi. Growth and Structural Transformation. In *Handbook of Economic Growth*, volume 2, chapter 6, pages 855–941. Elsevier, December 2014.
- Piyabha Kongsamut, Sergio Rebelo, and Danyang Xie. Beyond Balanced Growth. *The Review of Economic Studies*, 68(4):pp. 869–882, 2001.

- Simon Kuznets. 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, 1957.
- David Lagakos and Michael E. Waugh. Selection, Agriculture, and Cross-Country Productivity Differences. *American Economic Review*, 103(2):948–80, 2013.
- A Maddison. Economic Growth and Structural Change in the Advanced Countries. In I. Leveson and J. W. Wheeler, editors, *Western Economies in Transition: Structural Change and Adjustment Policies in Industrial Countries*, pages 41–65. London: Croom Helm, 1980.
- Alan Manning. 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, November 2004.
- Kiminori Matsuyama. Increasing Returns, Industrialization, and Indeterminacy of Equilibrium. *The Quarterly Journal of Economics*, 106(2):pp. 617–650, 1991.
- Francesca Mazzolari and Giuseppe Ragusa. Spillovers from High-skill Consumption to Low-skill Labor Markets. *The Review of Economics and Statistics*, 95(1):74 – 86, March 2013.
- Peter B. Meyer and Anastasiya M. Osborne. Proposed Category System for 1960-2000 Census Occupations. Working Papers 383, U.S. Bureau of Labor Statistics, September 2005.
- Guy Michaels, Ashwini Natraj, and John Van Reenen. Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 years. *The Review of Economics and Statistics*, 96(1):60 – 77, March 2014.
- L. Rachel Ngai and Barbara Petrongolo. Gender Gaps and the Rise of the Service Economy. Discussion Paper 9970, Centre for Economic Policy Research, May 2014.
- L. Rachel Ngai and Christopher A. Pissarides. Structural Change in a Multisector Model of Growth. *The American Economic Review*, 97(1):pp. 429–443, 2007.

L. Rachel Ngai and Christopher A. Pissarides. Trends in Hours and Economic Growth. *Review of Economic Dynamics*, 11(2):239 – 256, 2008.

Steven Ruggles, J. Trent Alexander, Katie Genadek, Ronald Goeken, Matthew B. Schroeder, and Matthew Sobek. Integrated Public Use Microdata Series. Technical report, Minnesota Population Center, Minneapolis, MN, 2010. Version 5.0.

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

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

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/industries or in the military. Our employment measure is the product of weeks worked times usual number of hours per week.⁴⁰ We compute hourly wages as earnings divided by the product of usual hours and weeks worked.

To construct the 30-year change graphs of Figure 1 and 10, and the 10-year change graphs of Figure 11 we follow the methodology used in Autor et al. (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. Figures 1, 10, and 11 plot the smoothed changes in average log hourly wages and total hours worked at each percentile of the occupational skill distribution. These skill percentiles are constructed by ranking the balanced occupations according to their 1950 (Figure 10 and top row of Figure 11) or 1980 mean hourly wages (Figure 1 and bottom row of Figure 11), and then splitting them into 100 groups, each making up 1 percentile of 1950 or 1980 employment.

Figure 10 shows the change in log real hourly wages and employment, similarly as Figure 1, with the difference that the ranking of occupations is based on their 1950 log

⁴⁰Since 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.

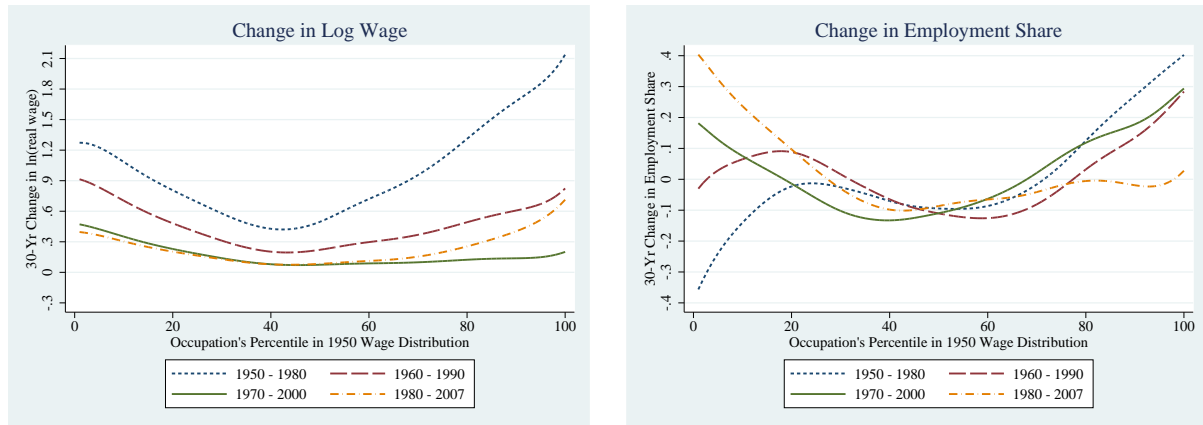


Figure 10: Smoothed wage and employment polarization 1950 ranking

Notes: Data and left and right panel same as in Figure 1, except occupations are ranked based on their 1950 mean wages.

real hourly wage. The graph reinforces the message of Figure 1. The left panel shows that wages have been polarizing from 1950 onwards, with the polarization most pronounced in the earlier periods. The right panel shows that the polarization of employment is present in all 30-year periods starting from 1960, with the most pronounced polarization between 1970-2000.

Figure 11 shows the wage and employment change of occupations for 10-year periods, with occupations ranked based on the 1950 wages (top row) and the 1980 wages (bottom row). These graphs show that polarization does not happen on a decade-by-decade basis. In some decades the top gains, while in others the bottom, but it is never the middle-wage occupations that gain the most in terms of wages or employment.

In the text we document polarization in terms of occupations for 183 and 10 occupation categories (in Figure 1 and Figure 2 respectively), here we show it for in an even coarser classification. As in Acemoglu and Autor (2011) we classify occupation groups into the following categories: manual, routine, and abstract.⁴¹ Figure 12 shows the patterns of polarization both in terms of wages and employment shares between 1950 and 2007 for these three broad categories. The right 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 com-

⁴¹Acemoglu and Autor (2011) have a similar graph of the path of employment shares of four occupation categories (abstract, routine cognitive, routine non-cognitive, manual) between 1960 and 2007. Here we show for 3 categories, starting from 1950, and more importantly, we also show the path of relative occupational wages.

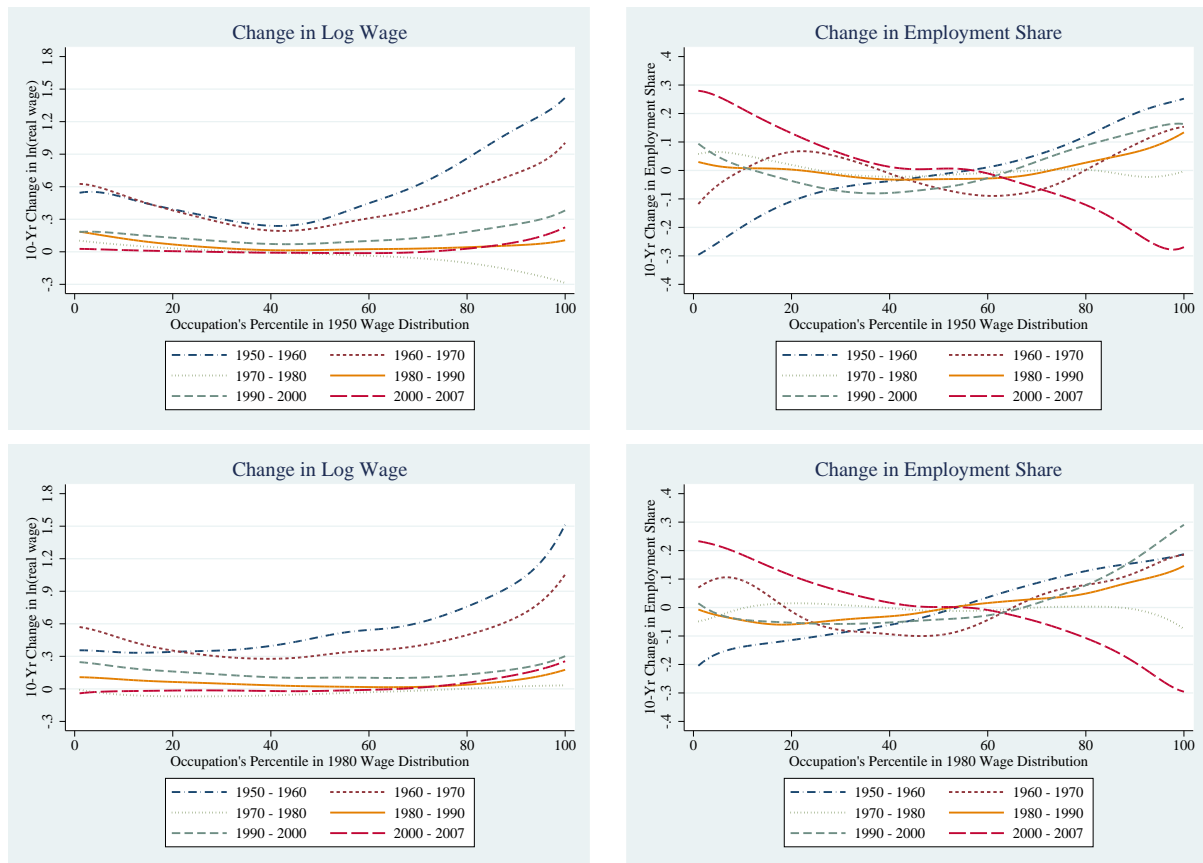


Figure 11: Smoothed wage and employment polarization, 10-year change

Notes: Data and left and right panel as in Figure 1. All panels show 10-year changes rather than 30-year changes. Occupations are ranked based on their 1950 mean wages in the top two panels, and based on their 1980 mean wages in the bottom two panels.

pression until 1960, has been steadily increasing. The left 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, also in terms of these three broad occupations there is clear evidence for polarization.

Categorization of occupations

Following Acemoglu and Autor (2011) we classify occupations into three categories, which are used in Figure 12:

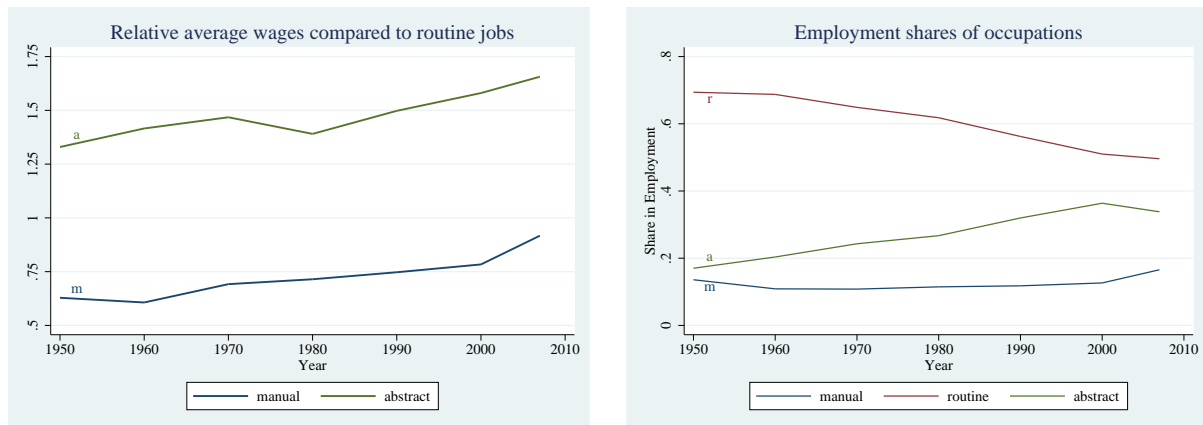


Figure 12: Polarization for broad occupations

Notes: Relative average wages and employment shares (in terms of hours) are calculated from the same data as in Figure 1. For details of the occupation classification see below.

- 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 3:

- Low-skilled services: personal services, entertainment, low-skilled transport (bus service and urban transit, taxicab service, trucking service, warehousing and storage, services incidental to transportation), low-skilled business and repair services (automotive rental and leasing, automobile parking and carwashes, automotive repair and related services, electrical repair shops, miscellaneous repair services), retail trade, wholesale trade;
- Manufacturing: mining, construction, manufacturing;
- High-skilled services: professional and related services, finance, insurance and real estate, communications, high-skilled business services (advertising, services to dwellings

and other buildings, personnel supply services, computer and data processing services, detective and protective services, business services not elsewhere classified), communications, utilities, high-skilled transport (railroads, U.S. Postal Service, water transportation, air transportation), public administration.

Table 5 summarizes the descriptive statistics for sectoral employment:

Table 5: Descriptive statistics by industry

	low-skilled services		manufacturing		high-skilled services	
	1960	2007	1960	2007	1960	2007
Highschool Dropout	55.15%	10.45%	56.29%	13.34%	30.68%	2.60%
Highschool Graduate	29.43%	36.77%	27.19%	39.98%	30.82%	19.19%
Some College	11.09%	33.50%	9.86%	26.79%	16.94%	30.84%
College Degree	3.82%	14.88%	5.63%	14.42%	14.26%	27.86%
Postgraduate	0.51%	4.40%	1.03%	5.48%	7.29%	19.50%
Avg Yrs of Education	10.26	13.08	10.21	12.86	12.21	14.69
Female Share	33.13%	48.00%	18.66%	21.85%	38.09%	54.35%
Foreign-Born Share	6.23%	18.05%	6.63%	20.02%	5.04%	12.88%

Occupation and sector premia

In Figures 3 and 12 as well as in our quantitative exercise we focused on relative average residual wages. We obtain these by regressing log hourly wages on sector dummies 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.

Table 6: Regression of log hourly wages: sector effects

Year	1950	1960	1970	1980	1990	2000	2007
low-sk. serv.	-0.28***	-0.31***	-0.22***	-0.19***	-0.20***	-0.17***	-0.18***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
high-sk. serv.	-0.03***	0.02***	0.08***	0.08***	0.14***	0.17***	0.21***
	(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	0.21	0.25	0.21	0.21	0.21	0.18	0.19

Standard errors in parentheses

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

Table 7: Regression of log hourly wages: occupation effects

Year	1950	1960	1970	1980	1990	2000	2007
manual	-0.41*** (0.01)	-0.42*** (0.00)	-0.33*** (0.00)	-0.28*** (0.00)	-0.24*** (0.00)	-0.19*** (0.00)	-0.10*** (0.00)
abstract	0.17*** (0.00)	0.27*** (0.00)	0.32*** (0.00)	0.31*** (0.00)	0.39*** (0.00)	0.44*** (0.00)	0.50*** (0.00)
Controls	yes	yes	yes	yes	yes	yes	yes
Observations	113635	459564	579290	958318	1094458	1235282	1308885
R^2	0.23	0.28	0.25	0.26	0.27	0.26	0.26

Standard errors in parentheses

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

Tables 6 and 7 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 sector 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 occupation dummies.

In the text we show the coefficients on the sectoral dummies from wage regressions, and in Figure 12 the relative average occupational wages. In Figure 13 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.

The role of gender and age composition changes

Figure 14 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 1960 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. This implies that the changing composition of the labor force is not

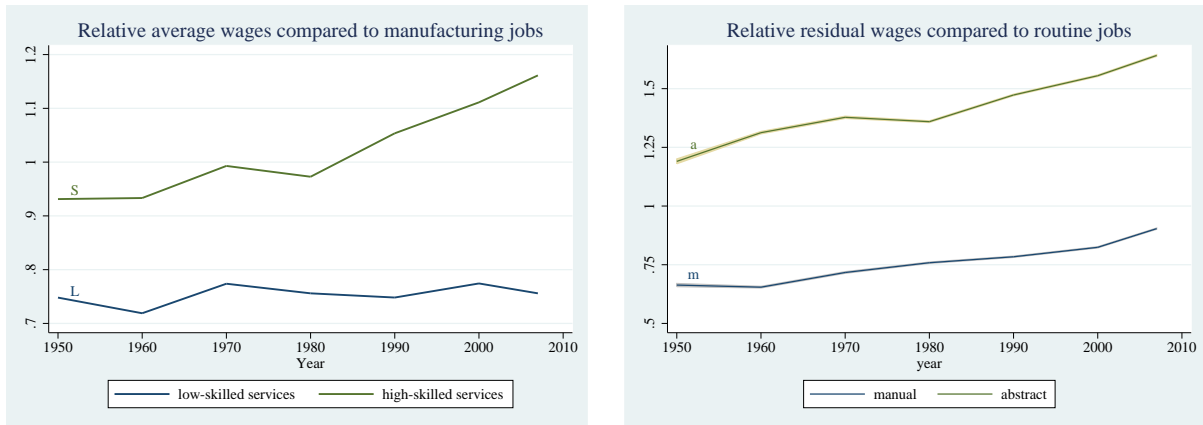


Figure 13: Wage polarization for sectors and occupations

Notes: Same data and classification as in Figure 3 and 12. 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 7.

the main driving force of the evolution of sectoral employment.

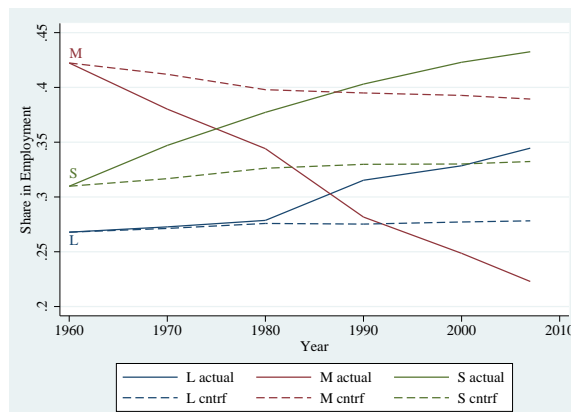


Figure 14: Counterfactual exercise: only changes in the gender-age composition of the labor force

Notes: Employment shares (in terms of hours) are calculated from the same data as in Figure 3. 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 Table 1 of the main text we showed a shift-share decomposition for the changes in occupational employment between 1950 and 2007, and alternatively between 1960 and 2007. In Table 8 we show this decomposition of employment share changes into a between-industry and a within-industry component for each decade. While we find a declining contribution of between-industry shifts since 1980, which might be due

routinization then taking off, again we find that a sizable part of the occupational employment share changes is due to shifts between industries.

Table 8: Decomposition of the changes in occupational employment shares by decade

	1950-60	1960-70	1970-80	1980-90	1990-00	2000-07
3 occupations, 3 sectors						
Manual						
Total Δ	-2.71	-0.07	0.67	0.31	0.85	3.93
Between Δ	-0.94	0.55	0.47	0.95	0.47	0.44
Within Δ	-1.76	-0.63	0.21	-0.65	0.39	3.48
Routine						
Total Δ	-0.65	-3.86	-3.09	-5.57	-5.24	-1.39
Between Δ	0.94	-1.41	-1.22	-1.58	-0.98	-0.70
Within Δ	-1.59	-2.45	-1.86	-3.99	-4.26	-0.69
Abstract						
Total Δ	3.35	3.93	2.41	5.27	4.39	-2.54
Between Δ	0.00	0.85	0.76	0.63	0.51	0.26
Within Δ	3.35	3.08	1.66	4.63	3.87	-2.80
10 occupations, 11 industries						
Manual						
Total Δ	-2.71	-0.07	0.67	0.31	0.85	3.93
Between Δ	-1.51	0.71	0.73	1.16	0.78	0.67
Within Δ	-1.19	-0.78	-0.06	-0.85	0.07	3.26
Routine						
Total Δ	-0.64	-3.86	-3.09	-5.57	-5.24	-1.39
Between Δ	0.85	-2.39	-1.96	-2.21	-1.80	-1.02
Within Δ	-1.49	-1.47	-1.12	-3.36	-3.44	-0.36
Abstract						
Total Δ	3.35	3.93	2.41	5.27	4.39	-2.54
Between Δ	0.67	1.69	1.23	1.05	1.02	0.36
Within Δ	2.69	2.25	1.18	4.21	3.37	-2.90

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 time interval given at the top. The top panel uses 3 occupations and 3 sectors, the bottom panel 10 occupations and 11 industries.

As an alternative way to assess 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 1960 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 15 shows the resulting time

series (dashed) and the actual data (solid). This exercise shows that if there had been only within-industry shifts, qualitatively the employment of the occupation categories would have evolved as in the actual data, but that quantitatively they cannot explain all of the changes. We therefore conclude that also between-industry shifts account for the polarization of occupational employment.

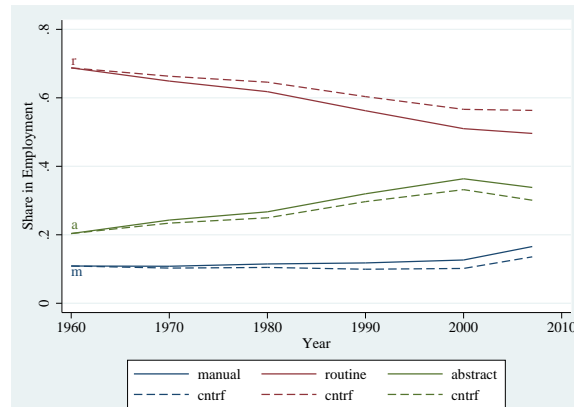


Figure 15: Counterfactual exercise: only-within industry shift of occupations

Notes: Employment shares (in terms of hours) are calculated from the same data as in Figure 13. 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.

Alternative shift-share decomposition

We also conduct an alternative shift-share decomposition, where we use industry level value added shares instead of employment shares. We construct hybrid occupational employment shares as

$$\tilde{E}_{ot} = \sum_i VA_{it} \lambda_{oit},$$

where VA_{it} is the share of industry i in total value added in period t , and λ_{oit} is the share of occupation o , industry i employment within industry i employment in period t , as defined earlier. In general $\tilde{E}_{ot} \neq E_{ot}$, where E_{ot} is simply the share of a occupation o in total employment, and which is equal to $E_{ot} = L_{ot}/L_t = \sum_i E_{it} \lambda_{oit}$, where E_{it} is the share of a industry i in total employment, as before.

Given these hybrid occupational employment shares, we can decompose their change into a part that is driven by within industry occupational employment share changes,

and a part that is driven by between industry shifts of value added.⁴²

$$\Delta \tilde{E}_{ot} = \underbrace{\sum_i \lambda_{oi} \Delta V A_{it}}_{\equiv \Delta \tilde{E}_{ot}^B} + \underbrace{\sum_i V A_i \Delta \lambda_{oit}}_{\equiv \Delta \tilde{E}_{ot}^W}.$$

Table 9: Decomposition of changes in hybrid occupational employment shares

Constructed employment shares		
1960-2007		
	3 x 3	10 x 11
Manual		
Total Δ	3.65	3.35
Between Δ	1.32	0.98
Within Δ	2.36	2.37
Routine		
Total Δ	-19.53	-18.50
Between Δ	-5.29	-6.46
Within Δ	-14.24	-12.50
Abstract		
Total Δ	16.09	15.36
Between Δ	4.18	5.69
Within Δ	11.90	9.68

Notes: Same occupational employment share data as in Figure 1. The value added industry data come from the BEA. For each occupational category, the first row presents the total change, the second the between-industry component, and the third the within-industry component over the period 1960-2007. The first column uses 3 occupations and 3 sectors, column two uses 10 occupations and 11 industries for the decomposition.

Table 9 shows the changes in our hybrid occupational employment shares and their decomposition between 1960 and 2007, into a between-industry and a within-industry component. The value added data comes from the Bureau of Economic Analysis (BEA). Due to the lack of value added data for finer industry categories before the 1960s, we decompose changes between 1960 and 2007. Table 9 suggests that between one fourth and one third of occupational employment changes are driven by between industry phenomena, regardless of whether we decompose 3 occupations in 3 sectors, or 10 occupations in 11 sectors. The importance of the between-industry component

⁴²The change driven by shifts between sectors is calculated as the weighted sum of the change in sector i 's value added share, $\Delta V A_{it}$, where the weights are the average employment 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 value added share of sector i , $V A_i = (V A_{it} + V A_{i0})/2$.

seems to be somewhat smaller than in the standard shift-share decomposition, but it is nonetheless quite a significant share of the overall change.

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:

$$\varepsilon \frac{d\widehat{a}_m}{\widehat{a}_m} + \frac{dL_l}{L_l} - \frac{dN_m}{N_m} = (1 - \varepsilon) \frac{d\frac{A_m}{A_l}}{\frac{A_m}{A_l}}, \quad (24)$$

$$\varepsilon \left(\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s} \right) + \frac{dN_s}{N_s} - \frac{dN_m}{N_m} = (1 - \varepsilon) \frac{d\frac{A_m}{A_s}}{\frac{A_m}{A_s}}. \quad (25)$$

Applying the Leibniz 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

$$dL_l(\widehat{a}_m, \widehat{a}_s) = \frac{\partial L_l}{\partial \widehat{a}_m} d\widehat{a}_m + \frac{\partial L_l}{\partial \widehat{a}_s} d\widehat{a}_s = \underbrace{\int_0^{\widehat{a}_s} f(\widehat{a}_m, a_s) da_s}_{\equiv C_1} \cdot d\widehat{a}_m + \underbrace{\int_0^{\widehat{a}_m} f(a_m, \widehat{a}_s) da_m}_{\equiv C_2} \cdot d\widehat{a}_s, \quad (26)$$

$$dN_m(\widehat{a}_m, \widehat{a}_s) = - \underbrace{\int_0^{\widehat{a}_s} f(\widehat{a}_m, a_s) da_s}_{=C_1} \cdot \widehat{a}_m d\widehat{a}_m - \underbrace{\int_{\widehat{a}_m}^{\infty} a_m^2 f(a_m, \frac{\widehat{a}_s}{\widehat{a}_m} a_m) da_m}_{\equiv C_3} \cdot \frac{\widehat{a}_s}{\widehat{a}_m} \left(\frac{d\widehat{a}_m}{\widehat{a}_m} - \frac{d\widehat{a}_s}{\widehat{a}_s} \right), \quad (27)$$

$$dN_s(\hat{a}_m, \hat{a}_s) = - \underbrace{\int_0^{\hat{a}_m} f(a_m, \hat{a}_s) da_m}_{=C_2} \cdot \hat{a}_s d\hat{a}_s + \underbrace{\int_{\hat{a}_s}^{\infty} a_s^2 f\left(\frac{\hat{a}_m}{\hat{a}_s} a_s, a_s\right) da_s}_{\equiv C_4} \cdot \frac{\hat{a}_m}{\hat{a}_s} \left(\frac{d\hat{a}_m}{\hat{a}_m} - \frac{d\hat{a}_s}{\hat{a}_s} \right), \quad (28)$$

similarly

$$dL_m(\hat{a}_m, \hat{a}_s) = - \underbrace{\int_0^{\hat{a}_s} f(\hat{a}_m, a_s) da_s}_{=C_1} \cdot d\hat{a}_m - \underbrace{\int_{\hat{a}_m}^{\infty} a_m f\left(a_m, \frac{\hat{a}_s}{\hat{a}_m} a_m\right) da_m}_{\equiv C_5} \cdot \frac{\hat{a}_s}{\hat{a}_m} \left(\frac{d\hat{a}_m}{\hat{a}_m} - \frac{d\hat{a}_s}{\hat{a}_s} \right), \quad (29)$$

$$dL_s(\hat{a}_m, \hat{a}_s) = - \underbrace{\int_0^{\hat{a}_m} f(a_m, \hat{a}_s) da_m}_{=C_2} \cdot d\hat{a}_s + \underbrace{\int_{\hat{a}_s}^{\infty} a_s f\left(\frac{\hat{a}_m}{\hat{a}_s} a_s, a_s\right) da_s}_{\equiv C_6} \cdot \frac{\hat{a}_m}{\hat{a}_s} \left(\frac{d\hat{a}_m}{\hat{a}_m} - \frac{d\hat{a}_s}{\hat{a}_s} \right). \quad (30)$$

Plugging these into (24) and (25) and re-arranging we get:

$$\begin{aligned} \frac{d\hat{a}_m}{\hat{a}_m} \underbrace{\left[\varepsilon + \frac{C_1 \hat{a}_m}{L_l} + \frac{C_1 (\hat{a}_m)^2}{N_m} + \frac{C_3 \frac{\hat{a}_s}{\hat{a}_m}}{N_m} \right]}_{\equiv B_1 > 0} + \frac{d\hat{a}_s}{\hat{a}_s} \underbrace{\left[\frac{C_2 \hat{a}_s}{L_l} - \frac{C_3 \frac{\hat{a}_s}{\hat{a}_m}}{N_m} \right]}_{\equiv B_2} &= \underbrace{(1 - \varepsilon) \frac{d \frac{A_m}{A_l}}{\frac{A_m}{A_l}}}_{\equiv D_1}, \\ \frac{d\hat{a}_m}{\hat{a}_m} \underbrace{\left[\varepsilon + \frac{C_4 \frac{\hat{a}_m}{\hat{a}_s}}{N_s} + \frac{C_1 (\hat{a}_m)^2}{N_m} + \frac{C_3 \frac{\hat{a}_s}{\hat{a}_m}}{N_m} \right]}_{\equiv B_3 > 0} - \frac{d\hat{a}_s}{\hat{a}_s} \underbrace{\left[\varepsilon + \frac{C_2 (\hat{a}_s)^2}{N_s} + \frac{C_4 \frac{\hat{a}_m}{\hat{a}_s}}{N_s} + \frac{C_3 \frac{\hat{a}_s}{\hat{a}_m}}{N_m} \right]}_{\equiv B_4 > 0} &= \underbrace{(1 - \varepsilon) \frac{d \frac{A_m}{A_s}}{\frac{A_m}{A_s}}}_{\equiv D_2}. \end{aligned}$$

This leads to

$$\begin{aligned} \frac{d\hat{a}_s}{\hat{a}_s} &= \frac{B_3 D_1 - B_1 D_2}{B_3 B_2 + B_1 B_4}, \\ \frac{d\hat{a}_m}{\hat{a}_m} &= \frac{D_2 B_2 + B_4 D_1}{B_3 B_2 + B_1 B_4}, \end{aligned}$$

where $B_3 B_2 + B_1 B_4 > 0$ can be easily verified by multiplying out the terms. Hence to determine the response in \hat{a}_m and in \hat{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:

$$\frac{d\hat{a}_s}{\hat{a}_s} = \frac{D}{B_3B_2 + B_1B_4}(B_3 - B_1) = \frac{D}{B_3B_2 + B_1B_4} \left(\frac{C_4 \frac{\hat{a}_m}{\hat{a}_s}}{N_s} - \frac{C_1 \hat{a}_m}{L_l} \right), \quad (31)$$

$$\frac{d\hat{a}_m}{\hat{a}_m} = \frac{D}{B_3B_2 + B_1B_4}(B_2 + B_4) = \frac{D}{B_3B_2 + B_1B_4} \left(\varepsilon + \frac{C_2(\hat{a}_s)^2}{N_s} + \frac{C_4 \frac{\hat{a}_m}{\hat{a}_s}}{N_s} + \frac{C_2 \hat{a}_s}{L_l} \right) > 0. \quad (32)$$

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

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

This together with (27) and (29) imply that N_m and L_m always decrease. These changes are:

$$\begin{aligned} dN_m(\hat{a}_m, \hat{a}_s) &= - \left(C_1(\hat{a}_m)^2 \frac{d\hat{a}_m}{\hat{a}_m} + C_3 \frac{\hat{a}_s}{\hat{a}_m} \left(\frac{d\hat{a}_m}{\hat{a}_m} - \frac{d\hat{a}_s}{\hat{a}_s} \right) \right) < 0, \\ dL_m(\hat{a}_m, \hat{a}_s) &= - \left(C_1 \hat{a}_m \frac{d\hat{a}_m}{\hat{a}_m} + C_5 \frac{\hat{a}_s}{\hat{a}_m} \left(\frac{d\hat{a}_m}{\hat{a}_m} - \frac{d\hat{a}_s}{\hat{a}_s} \right) \right) < 0. \end{aligned}$$

By plugging in (31) and (32) into (26) we can show that employment in sector L increases:

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

By plugging in (8) and (7) into (28) we can show that effective employment in sector S increases:

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

□

Quantitative results

Normalization of the mean of a_m and a_s

In what follows we show that the calibration and the transition path of the model when setting the mean of a_m and a_s to 1 is exactly the same as when setting the mean of a_m to c and the mean of a_s to d . As explained in the main text, we calibrate the model in two steps, first (for a given normalization for the means), we identify the minimum labor efficiencies in sector M (\underline{a}_m) and S (\underline{a}_s), which allow the employment shares and relative average wages to be matched at the same time. In this first step we also identify the cutoff efficiencies \hat{a}_m and \hat{a}_s , which given this distribution would give rise to the targeted employment shares and relative average wages. In the second step, we calibrate τ_l and τ_s for which, given the distribution, the equilibrium cutoffs are the ones found in the previous step.

Let $\underline{a}_m^{1,1}, \underline{a}_s^{1,1}, \tau_l^{1,1}$ and $\tau_s^{1,1}$ denote the calibrated parameter values when both means are normalized to 1, and the equilibrium cutoff efficiencies be denoted by $\hat{a}_m^{1,1}$ and $\hat{a}_s^{1,1}$. It can be shown that the calibration in case the means of a_m and a_s are normalized to c and d leads to $\underline{a}_m^{c,d} = c\underline{a}_m^{1,1}$, $\underline{a}_s^{c,d} = d\underline{a}_s^{1,1}$, $\tau_l^{c,d} = c^{\varepsilon-1}\tau_l^{1,1}$ and $\tau_s^{c,d} = (c/d)^{\varepsilon-1}\tau_s^{1,1}$, with equilibrium cutoffs given by $\hat{a}_m^{c,d} = c\hat{a}_m^{1,1}$ and $\hat{a}_s^{c,d} = d\hat{a}_s^{1,1}$. The calibration procedure thus leads to exactly the same initial equilibrium whether the means are normalized to $(1, 1)$ or (c, d) .

What is left to show is that these two different setups lead to the same equilibrium changes as relative productivities change. Let us denote the proportional change in $\hat{a}_m^{1,1}$ by u and the proportional change in $\hat{a}_s^{1,1}$ by v following a change in relative productivities A_m/A_l and A_m/A_s . It can be shown that cutoffs $u\hat{a}_m^{c,d}$ and $v\hat{a}_s^{c,d}$ clear the market for the new relative productivities, and therefore lead to the same change in employment shares and equilibrium wages as in the case when the means are normalized to $(1, 1)$.

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 et al. (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 (1960-2007) is 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 labor efficiency distribution which we have calibrated. Given this distribution we can calculate the sector-of-work cutoffs \hat{a}_m and \hat{a}_s needed to match the employment shares calculated from the Census and ACS data. Given the cutoffs and the distribution of labor efficiencies, we can calculate the average efficiency in each sector for these years (1960, 1970, 1980, 1990, 2000, 2007). We then generate effective employment from these average labor efficiencies and the raw employment data from the BEA. The adjusted labor productivity is just the sectoral value added divided by the effective employment.

Employment and wage dynamics with decennial and adjusted labor productivity

Figure 16 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 4). As manufacturing productivity growth accelerates relative to services in the 1980, also when adjusting for workers' self-selection, this figure is very similar to Figure 8 of the main text.

Robustness

In section 3.6 of the main text we summarized how our result change when assuming a different underlying distribution of sectoral efficiencies and when varying the elas-

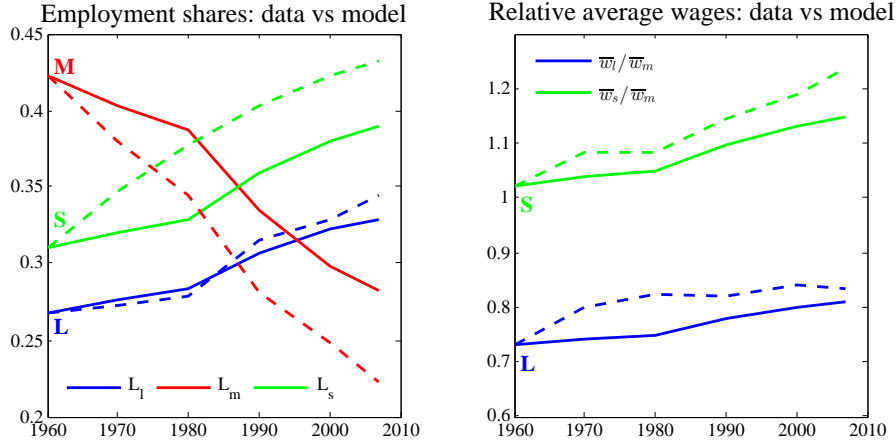


Figure 16: Transition of the model with selection-adjusted decennial productivity growth rates

ticity of substitution between goods and services (measured in value-added terms).

Table 10: Robustness checks: different distributions vs the data

distribution	$corr(a_m, a_s)$	$Var(a_m)$	$Var(a_s)$	employment share Δ			rel. avg. wage Δ	
				L	M	S	L to M	S to M
uniform	0	0.12	0.32	6.92	-16.10	9.18	12.28	14.58
lognormal	0	0.106	0.143	6.97	-17.09	10.12	5.48	5.51
	0.3	0.095	0.122	7.06	-16.67	9.60	4.94	3.88
	0.5	0.089	0.110	7.12	-16.35	9.23	4.59	2.90
tr. normal	0	0.143	0.229	6.93	-16.70	9.77	8.31	9.51
	0.3	0.122	0.184	7.04	-16.24	9.20	7.62	7.03
	0.5	0.110	0.156	7.11	-15.93	8.82	7.18	5.60
data				7.67	-19.94	12.27	13.97	21.16

Notes: The top row shows our baseline calibration of the uniform distribution, the next three rows a lognormal distribution with different correlations between a_m and a_s , and the last three rows truncated (at zero) normal distributions. Column 2 shows the correlation coefficient of the distribution, 3 and 4 show the variance of a_m and a_s implied by the calibration, 5 to 7 show the percentage point change in employment shares, while the last two show the percentage change in relative average wages. The last row contains the change in these same measures between 1960 and 2007 in the data.

Table 10 shows the model's predictions for three calibrations of the lognormal and the normal distribution truncated at zero for both a_m and a_s . For both class of distributions the mean of a_m and a_s can be normalized to one without loss of generality. The variance-covariance matrix of the bivariate distribution as well as τ_l and τ_s , altogether 5 parameters, need to be calibrated. Since we only have 4 moments in the data to match, we calibrate both the lognormal and the truncated normal distribution for

three different values of correlation between a_m and a_s : 0, 0.3 and 0.5. The qualitative predictions of the model are the same for all distributions considered. The quantitative results are virtually the same for employment shares, while the predictions for relative average wages vary more with the distribution. As discussed in the main text, this is driven by the differences in the correlation and calibrated variances.

Table 11: Robustness checks: different model specifications vs the data

ε	τ_l	τ_s	productivity	employment share Δ			relative average wage Δ	
				L	M	S	L to M	S to M
0.002	0.49	0.91	raw	6.92	-16.10	9.18	12.28	14.58
			adjusted	6.05	-13.97	7.92	10.60	12.38
0.02	0.49	0.91	raw	6.81	-15.77	8.97	12.03	14.21
			adjusted	5.94	-13.67	7.73	10.38	12.07
0.2	0.48	0.88	raw	5.57	-12.52	6.95	9.53	10.83
			adjusted	4.82	-10.79	5.97	8.19	9.21
data				7.67	-19.94	12.27	13.97	21.16

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 4). The last row contains the change in these same measures between 1960 and 2007 in the data.

In Table 11 we report the detailed results for the model's sensitivity with respect to ε , which we vary from the baseline value 0.002, as estimated by Herrendorf et al. (2013), to higher values in the range of plausible values reported by Ngai and Pissarides (2008). For each value of ε , we recalibrate the other parameters to match moments of the 1960 data, as in section 3.4, and show the model's predictions over 1960-2007, both when feeding in the raw average productivity growth rates (as in the baseline simulation of Figure 6) and for the selection-adjusted growth rates (as in Figure 7). Qualitatively the patterns of the model-implied behavior of employment shares and relative average wages do not change with the elasticity of substitution. Quantitatively, the predictions naturally are affected, but even with the least favorable calibration (high ε and labor productivity growth adjusted for selection), the model predicts a substantial fraction of the sectoral wage and employment patterns of the last 50 years.