Job Polarization and Structural Change*

Zsófia L. Bárány† Christian Siegel‡

27 August 2014

Abstract

Job polarization is a widely documented phenomenon in developed countries since the 1980s: employment has been shifting from middle to low- and high-income workers, while average wage growth has been slower for middle-income workers than at both extremes. We document 1) that polarization has started as early as the 1950s in the US, and 2) that this process is closely linked to the shift from manufacturing to services. Based on these observations we propose a structural change driven explanation for polarization. Productivity growth through raising national income leads to a disproportionate increase in the demand for high-end (luxury) services. To attract more workers into the high-skilled services, the wages in this sector have to grow at a faster pace than in the middle. The growing income of the wealthier part of the population in turn increases their demand for low-skilled services, leading to a partial marketization of home production, and a faster growth of the low-skilled service wages.

JEL codes: E24, J22, O41

Keywords: Job Polarization, Structural Change, Home Production

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

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

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

1 Introduction

The polarization of the labor market in terms of occupations is a widely documented phenomenon in the US and several European countries since the 1980s (Autor, Katz, and Kearney (2006), Goos and Manning (2007), and Goos, Manning, and Salomons (2009)). This phenomenon, besides the relative growth of wages and employment of high earning occupations, also entails the relative growth of wages and employment of low earning occupations. The leading explanation for polarization is the routinization hypothesis, which relies on the assumption that information and computer technologies (ICT) substitute for middle-skill and hence middle-earnings (routine) occupations, whereas they complement the high-skilled and high-earnings (abstract) occupations (Autor, Levy, and Murnane (2003), Autor, Katz, and Kearney (2006), Michaels, Natraj, and Van Reenen (2010), Goos, Manning, and Salomons (2011), Autor and Dorn (2012)).

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, is not the sole 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 Census data for the period 1950-2000 and American Community Survey (ACS) data for 2007 reveals some novel facts. First, while most of the literature on polarization focuses on occupations, we document that labor market polarization is present also in terms of broadly defined sectors: low-skilled and high-skilled service workers, who are at opposite ends of the earnings distribution have been gaining in terms of wages and employment at the expense of manufacturing workers. Second, we show that the loss in routinizable occupations is not uniform across sectors: routinizable employment only declined in the manufacturing sector. Third, we find that polarization has started as early as the 1950-1960s in the US. This implies that polarization started long before ICT or increased trade flows could have impacted the labor market. Observing a) that polarization seems to be a long-run phenomenon, b) that the middle earning jobs are in manufacturing, c) that manufacturing employment started to fall, while service employment started to increase in the 1950s-1960s, it is natural to investigate whether the structural shift of the economy is driving the polarization of the labor market.

Based on these observations we propose a structural change driven explanation for the joint polarization of wages and employment. In our general equilibrium model workers with heterogeneous ability select which sector to work in. As technology improves, the demand for both high- and low-skilled services increases disproportionately, which leads to a reallocation of labor from middle-earning manufacturing jobs to high-earning and low-earning service jobs. However, to attract more workers into these two sectors, the wages in these two sectors have to grow at a faster pace than in the middle. Finally, we calibrate the model and quantitatively assess the contribution of structural change – driven by both non-homothetic preferences and unbalanced technological progress – to the polarization of wages and employment.

In the model, there are three types of consumption goods: low-end service, manufacturing and high-

end service goods. Preferences over these three goods are non-homothetic: low-end services are necessities, while high-end services are luxury goods. Low-end services can be home produced or bought on the market, and as income increases a smaller fraction is home produced. Manufacturing goods and high-end services can only be market produced, but high-end services are only demanded at high enough income levels. As technology progresses in manufacturing and in the high-skilled service sector, the employment and production structure of the economy changes: as households gradually become wealthier, the demand for high-skilled services rises over-proportionally due to the non-homotheticity of preferences. This disproportionate demand rise puts an upward pressure on prices and wages in the high-skilled service sector. Consequently, more people sort into working in high-end service jobs. As wages increase, households decide to outsource a larger fraction of housework. This puts an upward pressure on wages at the bottom-end of the distribution, increasing the supply of these workers. Due to the non-homotheticity of preferences, the timing of the expansion of the high- and the low-skilled service sector can be different. While the non-homotheticity is relatively important, the desire to expand the consumption of high-skilled services dominates that of low-skilled services. This can be a potential explanation of the low-skilled service expansion becoming more pronounced later on in the data.

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

The structural change literature has documented for several countries that as income increases resources are shifted away from agriculture and from manufacturing towards services (Kuznets (1957), Maddison (1980)). In particular the employment share of manufacturing has been declining since the 1950s, while the employment share of services has been increasing. The literature has identified two economic forces that lead to structural transformation: preferences and technology. The preferences explanation relies on changes in aggregate income, which if preferences for the output of different sectors are not homothetic lead to a reallocation of resources across sectors (Caselli and Coleman II. (2001), Kongsamut et al (2001)). The technology explanation assumes that productivity growth is different across sectors, which with regular preferences leads to a shift of labor into the lower growth sector (Ngai and Pissarides (2007), Acemoglu and Guerrieri (2008)). The consensus seems to be that both mechanisms together are needed to explain the patterns observed in the data (Buera and Kaboski (2009), Herrendorf et al (2013), Boppart (2011)). Several papers have also established the importance of home production for structural transformation (Ngai and Pissarides (2008), Buera and Kaboski (2009), (2012a), (2012b)). In our model we rely on both mechanisms, and incorporate home production.

We extend the structural change literature in two ways. First, we allow heterogeneous workers to endogenously sort into different sectors. In the presence of differential productivity or demand growth, the optimal sector choices naturally change. Moreover, as the sorting into sectors change, the relative wages and prices are affected. Therefore, we can analyze the effects of structural change on relative

sectoral wages, which is not usual in models of structural change.[1]
Second, based on the job polarization phenomenon we distinguish between two types of services: low- and high-skilled. This is an important distinction, due to both the way they enter the utility of agents and the way they are produced. We believe that consumers enjoy these services in different ways. We model low-skilled services as substitutes for household production. We assume that there is a subsistence level of household production. On the other hand, we assume that high-skilled services are luxury goods, so the demand for them increases more than proportionately with income. In terms of production, high-skilled services can only be market produced, while low-skilled services can be home produced equally well.

Two popular explanations suggested for the polarization of occupations are the routinization hypothesis, and the consumption hypothesis. The routinization hypothesis relies on the assumption that information and computer technologies (ICT) substitute for middle-skill and hence middle-earnings (routine) occupations, whereas they complement the high-skilled and high-earnings (abstract) occupations (Autor, Levy, and Murnane (2003), Autor, Katz, and Kearney (2006), Michaels, Natraj, and Van Reenen (2010), Goos, Manning, and Salomons (2011), Autor and Dorn (2012)). It has been argued that much of the expansion of low-skilled occupations is driven by the expansion of low-skilled service jobs (Autor and Dorn (2012)). The routinization hypothesis suggests, that it is ICT that leads to the substitution of routine workers by machines, and which complements abstract workers. The displaced routine workers find either abstract or manual jobs, increasing the employment share of these occupations (Autor, Levy, and Murnane (2003), Autor, Katz, and Kearney (2006), Michaels, Natraj, and Van Reenen (2010), Goos, Manning, and Salomons (2011)). The routinization hypothesis, linked to ICT, is potentially a convincing explanation for the employment share patterns after the mid-1980s, but without any assumption on demands, it does not provide a mechanism through which the relative wage of manual workers compared to routine workers can increase. It also cannot provide an explanation for the patterns observed before ICT could have taken effect. The consumption spill-over argument, on the other hand, 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 (2007)). Our modeling of low-skilled service jobs as substitutes for home production is reminiscent of this argument.

2 Polarization in the data

In the empirical literature, polarization is mostly represented in terms of occupations. Following the methodology used in Autor, Katz, and Kearney (2006), Acemoglu and Autor (2011), and Autor and Dorn (2012), we plot the smoothed changes in employment shares and log real wages for a balanced panel of occupation categories ranked according to their 1950 and 1980 mean wages. The novelty in

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

these graphs is that we show these patterns going back until 1950, whereas most analyses look at data from only 1980 onwards. The top row of Figure 1 shows that there has been polarization in terms of real

Figure 1: Wage and employment polarization

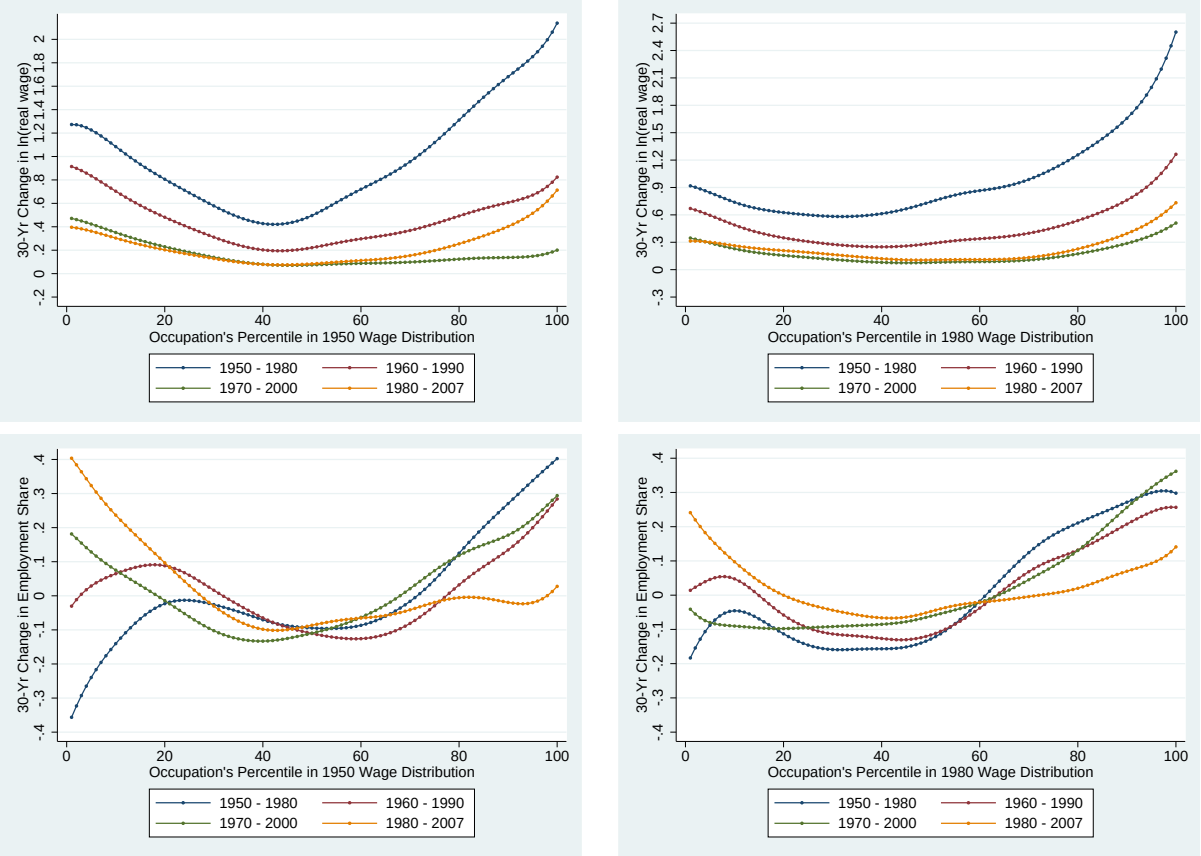

Notes: The data is taken from IPUMS US Census data for 1950, 1960, 1970, 1980, 1990, 2000 and American Community Survey (ACS) for 2007. The sample excludes agricultural occupations/industries and observations with missing wage data; the details are given in the appendix. Balanced occupation categories (185 of them) were defined by the authors based on Meyer and Osborne (2005) and Autor and Dorn (2012). The bottom two panels show the 30-year change in employment shares (calculated as hours supplied rather than persons), and the top two panels show the 30-year change in log hourly real wages (again labor supply weighted). In the left panels occupations are ranked based on their 1950 average wage, whereas in the right panels they are ranked according to their 1980 average wage.

wages in all 30-year periods. This polarization is present whether the occupations are ranked according to their 1950 or 1980 mean wages. The polarization of real wages is most pronounced in the first two 30-year intervals, but it is clearly discernible in the following ones as well from the slight U-shape of the smoothed changes. The picture is more mixed in terms of employment polarization (the bottom row of Figure 1): employment polarization is most pronounced in the last 30 years (1980-2007), but it seems to be present even in the earlier decades.[2]

These graphs – in line with the literature – plot the change in raw employment shares and in raw log real hourly wages. These changes also include the potential effects of the changing gender, age and race composition of the labor force.[3] These graphs also do not directly relate to the explanations put

[2]This does not necessarily hold for decade-by-decade analysis. Typically in some decades the top gains, whereas in others the bottom gains, but it is never the middle that grows the most in terms of employment shares. See graphs in Appendix.

[3]In figure 8 of the Appendix we show that the changing gender and age composition of the labor force does not account for

forward in the literature, as they show the employment share and wage changes for occupations ranked according to their mean wage, not based on their routinizability. Therefore we classify the occupation groups into the following categories: manual, routine, and abstract (as in Acemoglu et al (2011)), and show the patterns for these three broad categories. We also classify industries into three categories: low-skilled services, which are substitutes for household production; manufacturing; and high-skilled services, which are luxury goods. Manufacturing industries are as in the structural change literature (Buera et al (2012a), Kongsamut et al (2001), Ngai et al (2008), Herrendorf et al (2013)). We classify industries to be low-skilled services if they can be viewed as substitutes for household production.[4]

Figure 2: Polarization for industries and occupations

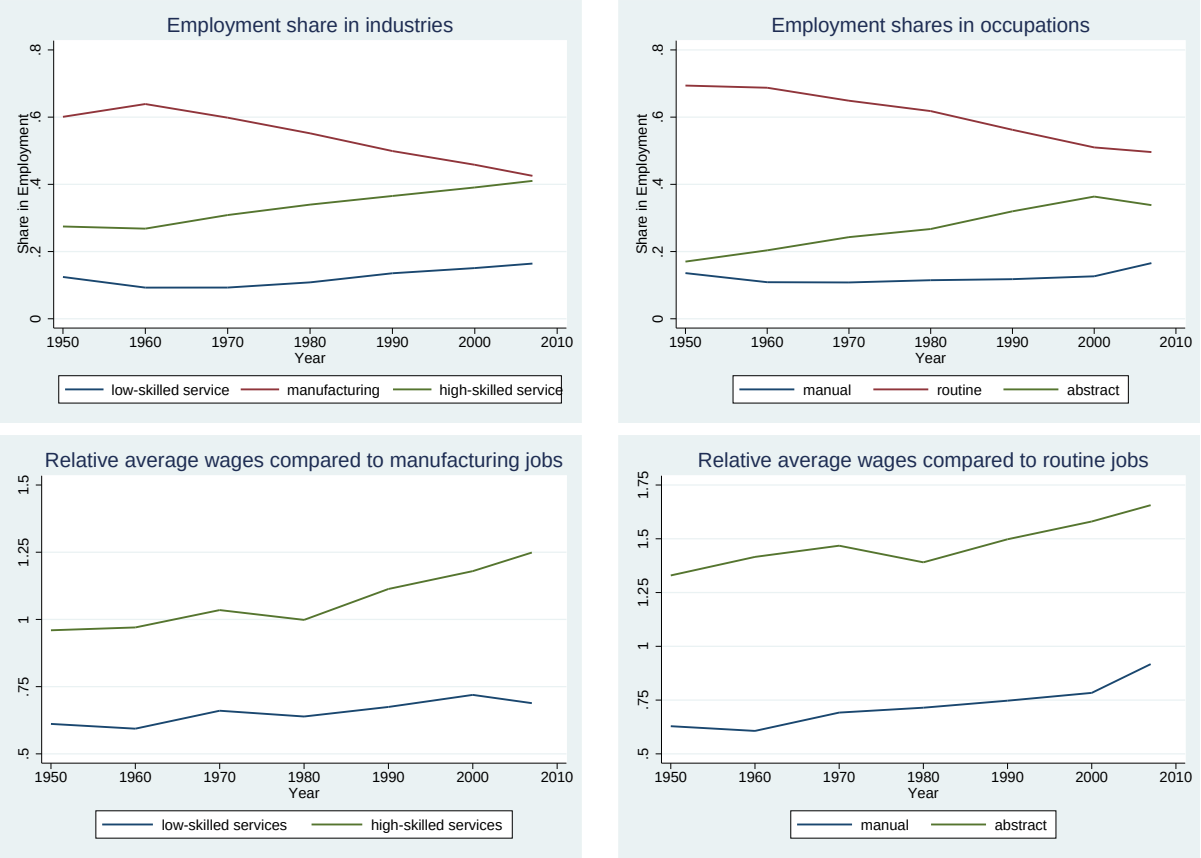

Notes: Employment shares (in terms of hours) and wages are calculated from the same data as in Figure 1. For details of the industry and occupation classification see text and the appendix.

Figure 2 shows the patterns of polarization both in terms of employment shares and wages for the above defined occupations and industries between 1950 and 2007. The patterns are strikingly similar between the graphs generated using industries (the left panels) and occupations (right panels). The

the observed employment polarization, albeit qualitatively working in that direction.

[4]Details on these classifications are given in the Appendix. While we are not using the educational attainment of employees in categorizing industries, the ranking that results from our distinction between substitutes for home production and luxurious services turns out to be such that indeed low-skilled services have a less skilled workforce than high-skilled services. The educational attainments of workers in low-skilled services and in manufacturing however is virtually identical; see Table 3.

top two panels show clear employment polarization both in terms of industries and occupations. The middle earning group (manufacturing/routine occupations) lost significantly in terms of employment share, the top (high-skilled services/abstract) gained, and the bottom (low-skilled services/manual) initially shrank, but then expanded. The bottom two panels show the change in average industry (or occupation) log hourly wage change in the given decade compared to the mean log hourly wage change in manufacturing (or routine occupations). In most decades average wages in low- and high-skilled services improved relative to manufacturing, while manual and abstract occupations also improved relative to routine occupations. Exceptions are the first and last decade, and the period between 1970-1980, when the high-skilled services and the abstract occupations lost. In this decade there was a secular compression in the skill premium, most probably due to forces outside the scope of this model.[5]

This striking similarity in the employment share and average wage path of the three broad industry and occupation classifications can be understood when considering the employment shares of the occupation categories in the three industry categories and vice versa. The top panel in Table 1 shows the employment shares in each industry-occupation cell. It can be seen that the largest numbers are on the diagonal, implying there is a tight correspondence between industries an occupations. The middle panel shows for each industry the fraction of employment coming from manual, routine and abstract occupations, while the bottom panel shows the opposite averaged between 1950-2007. The majority of workers in low-skilled services are in manual occupations, the majority in manufacturing are in routine occupations, and the majority in high-skilled services are in abstract occupations. The opposite is also true: the majority of routine workers are employed in manufacturing, and the majority of abstract workers are employed in the high-skilled service industry.

Table 1: Overlap in employment between industry and occupation

	low-skilled service	manufacturing	high-skilled service	total
manual	6.56	1.37	5.00	12.92
routine	3.92	40.22	12.45	56.60
abstract	2.86	8.90	18.73	30.48
total	13.34	50.48	36.18	100.00
manual	49.15	2.71	13.81	-
routine	29.40	79.67	34.42	-
abstract	21.45	17.62	51.76	-
total	100	100	100	-
manual	50.76	10.58	38.67	100
routine	6.93	71.07	22.00	100
abstract	9.39	29.18	61.43	100

Notes: Employment shares (in terms of hours) as percents are calculated from the same data as in Figure 1. Industry and occupation classification same as in Figure 2. The top panel shows the employment shares in each of the occupation-industry cells. The middle panel shows within each industry the employment share of different occupations, while the bottom panel shows within each occupation group the employment share of different industries.

It is informative to look at the changes in the employment shares in the various industry-occupation cells. Figure 3 shows that the employment share only declined in the routine-manufacturing cell,

[5]These forces include for instance the entrance of the baby boomers cohort into the labor market, who had an unusually high college enrollment rate probably in order to postpone the military draft.

Figure 3: The path of employment shares for industry-occupation cells

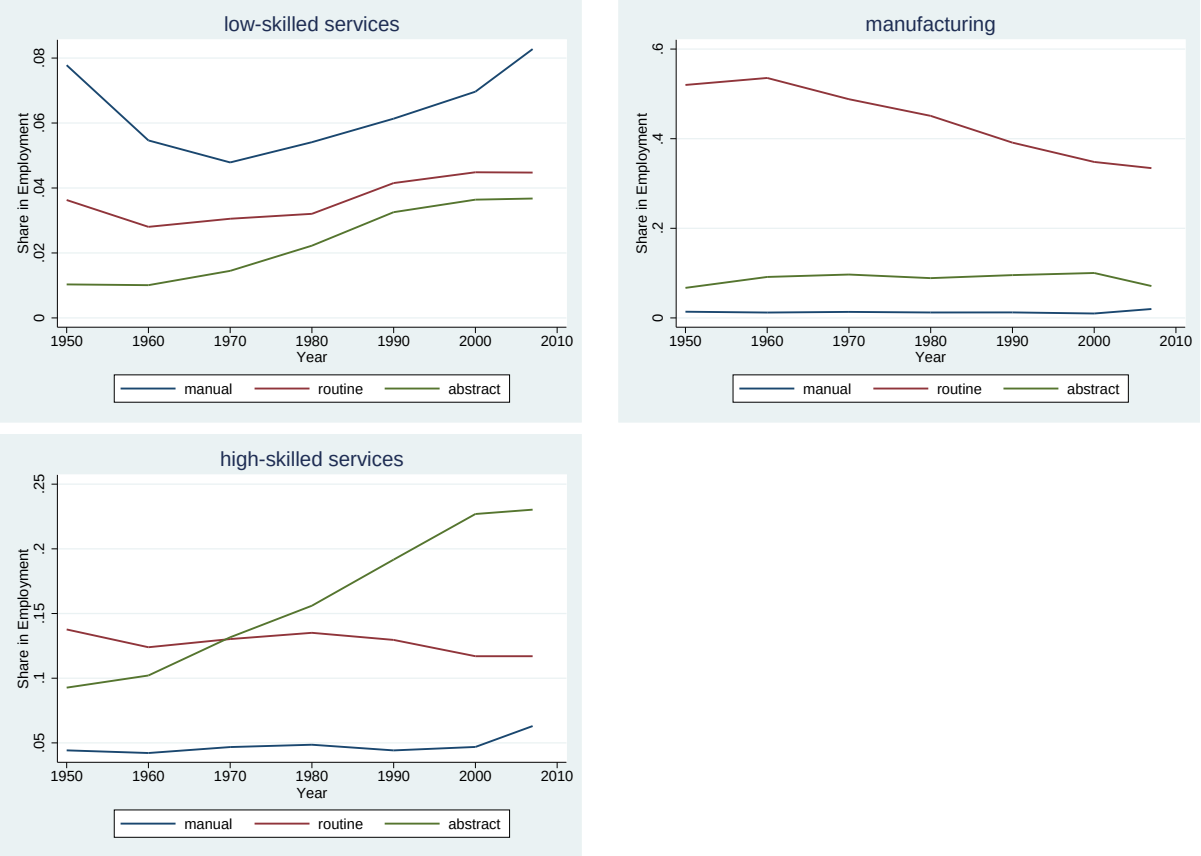

Notes: Employment shares (in terms of hours) are calculated from the same data as in Figure 1. Industry and occupation classification same as in Figure 2 and Table 1.

whereas the routine-high-skilled cell's employment share was stable, while the routine-low-skilled cell's employment share increased between 1950-2007. Therefore it seems that the decline in routinizable occupations is intrinsically linked to the decline in the manufacturing sector.

To further asses the importance of these employment reallocations between industries for the shifts in the broad occupation categories, we conduct the following counterfactual exercise: We fix the industry shares in employment (in terms of hours worked) at their 1950 levels and let the within-industry share of occupations follow their actual path, in order to compute how the occupational shares would have evolved in the absence of between industry shifts. Figure 4 shows the resulting time series (dashed) and the actual data (solid). The exercise shows that if there were only within-industry shifts qualitatively the employment of the occupation categories would have evolved as in the actual data, but the exercise also points out that quantitatively they cannot explain all of the reallocation. We therefore conclude that also between industry shifts account for the polarization of occupational employment. Since polarization occurs not in only in terms of occupations, but is also present in terms of broadly defined industries, starting as early as 1950-1960, we propose an explanation based on structural change.

Figure 4: Counterfactual exercise: only-within industry shift of occupations

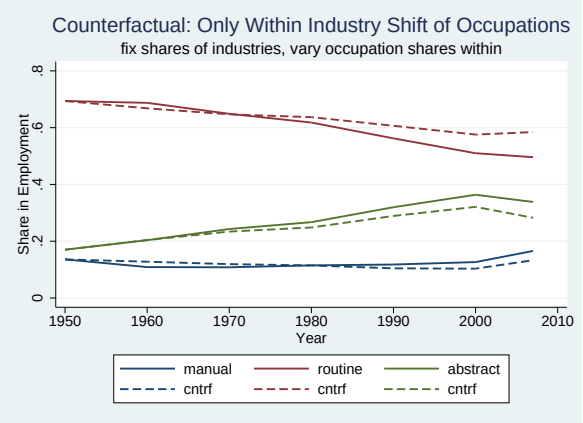

Notes: Employment shares (in terms of hours) are calculated from the same data as in Figure 2 and Table 1. The actual data is shown as solid lines, while the dashed line show how the occupational employment shares would have evolved in the absence of reallocations across industries.

3 Model

In order to illustrate the mechanism that is driving polarization, we present a parsimonious static model, and analyze its behavior as productivity levels increase across sectors. In this analysis we allow for both the non-homotheticity of preferences and uneven productivity growth across sectors. One novelty in our model is the distinction between low- and high-skilled services. We assume that low-skilled services are necessities, and can be equally well produced at home or on the market. High-skilled services, on the other hand, are luxury goods and can only be produced on the market. We also model the sorting of individuals into sectors based on the return to their labor, and thus are able to analyze the impact of structural change on the evolution of sectoral wage differentials.

We assume that the economy is populated by heterogenous agents, who all make individually optimal decisions about their sector of work. Furthermore these individuals are organized into a stand-in household, which collects all earnings, and provides the same consumption basket for each of its members. This household maximizes the household's utility subject to the household level budget constraint. Households derive utility from consuming high- and low-skilled services and manufacturing goods. Low-skilled services, which can be produced at home at a utility loss, or can be bought on the market from low-skilled service workers, are necessities: there is a subsistence amount required by each household. High-skilled services are luxury products in the sense that as income increases households spend an increasing fraction of their total consumption budget on these services.

Every individual chooses their sector of work to maximize earnings. We use a Roy-model type setup for the sector of work decision. We assume that in low-skilled services – which are substitutes for home production – 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 higher.

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

3.1 Sectors and production

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

The technology to produce high-skilled services is:

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

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

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

The technology to produce manufacturing goods is:

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

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

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

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

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

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

where A_l is productivity, and L_l is the raw units of labor (total amount of people) working in the low-

skilled service sector. Since sector L firms are also price takes, the wage per unit of raw labor in sector L has to satisfy:

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

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

3.2 Labor supply and demand for goods

The stand-in household consists of a measure one continuum of different types of members. Each member chooses which one of the three market sectors to supply his one unit of raw labor in. The household collects the earnings of all its members and decides how much low-skilled services, manufacturing goods and high-skilled services to buy on the market. Additionally it assigns each household member to spend the same amount of their leisure time in providing home production, which is a substitute for low-skilled services to the household. This implies that each household member will enjoy the same utility, independent of their innate differences.

3.2.1 Sector of work

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

Individuals are heterogeneous in their innate ability, $(a_1, a_2) \in \mathbf{R}^2_+$, which is drawn from a time invariant distribution $f(a_1, a_2)$. For simplicity we assume that a_1 denotes the individual's efficiency units of labor in manufacturing, while a_2 denotes his efficiency units in high-skilled services.[6] We further assume that each individual is equally productive when working in low-skilled services, as that only requires an individual's raw labor. Therefore the earnings on an individual with innate ability (a_1, a_2) in sector L is w_l, in sector M it is $a_1 w_m$, while if working in sector S it is $a_2 w_s$.

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

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

[6]The qualitative results of the model are unchanged if we assume that an individual's efficiency units of labor in manufacturing is $\alpha_m a_1 + (1 - \alpha_m)a_2$, while in high-skilled services it is $\alpha_s a_1 + (1 - \alpha_s)a_2$, and $\alpha_m \neq \alpha_s$.

cutoff values:

$$a_{1lm} \equiv \frac{w_l}{w_m} \tag{7}$$

$$a_{2ls} \equiv \frac{w_l}{w_s} \tag{8}$$

$$a_{2ms} \equiv \frac{w_m}{w_s}. \tag{9}$$

It is optimal for an individual with innate ability (a_1, a_2) *to work in sector* L *if and only if*

$$a_1 \leq a_{1lm} \quad \textit{and} \quad a_2 \leq a_{2ls}. \tag{10}$$

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

$$a_1 \geq a_{1lm} \quad \textit{and} \quad a_2 \leq a_{2ms} a_1. \tag{11}$$

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

$$a_2 \geq \max(a_{2ls}, a_{2ms} a_1). \tag{12}$$

Figure 5 shows this endogenous sorting behavior assuming that $a_1 \in [\underline{a}_1, \overline{a}_1]$ and $a_2 \in [\underline{a}_2, \overline{a}_2]$. Individuals who have low efficiency units or low productivity in both manufacturing and high-skilled services sort into low-skilled services (the blue area). Individuals with high enough manufacturing efficiency and relative to this a low high-skilled service efficiency sort into manufacturing jobs (the red area). While those individuals who have a high enough high-skilled service efficiency and relative to this a low manufacturing efficiency choose to work high-skilled services (the green area).

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

$$L_l = \int_{\underline{a}_1}^{a_{1lm}} \int_{\underline{a}_2}^{a_{2ls}} f(a_1, a_2) da_2 da_1 \tag{13}$$

$$N_m = \int_{a_{1lm}}^{\overline{a}_1} \int_{\underline{a}_2}^{a_{2ms} a_1} a_1 f(a_1, a_2) da_2 da_1 \tag{14}$$

$$N_s = \int_{\underline{a}_1}^{\overline{a}_1} \int_{\max(a_{2ls}, a_{2ms} a_1)}^{\overline{a}_2} a_2 f(a_1, a_2) da_2 da_1 \tag{15}$$

3.2.2 Demand for consumption goods and services

Household members derive utility from low-skilled services, manufacturing goods and high-skilled services. While manufacturing goods, c_m, and high-skilled services, c_s, can only be bought on the market, low-skilled services can be home produced, c_{lh}, and bought on the market, c_{lm}. The total

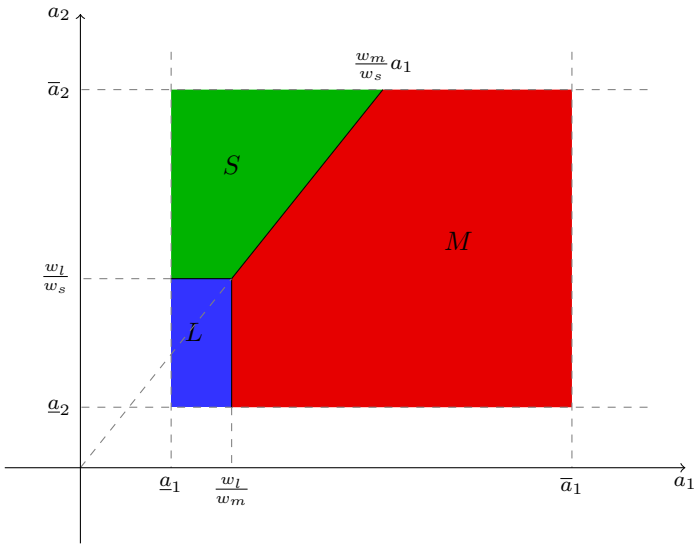

Figure 5: Optimal sector of work

Assuming that $a_1 \in [\underline{a}_1, \overline{a}_1]$ and $a_2 \in [\underline{a}_2, \overline{a}_2]$, the figure depicts the optimal sector of work choices given unit wage rates w_l, w_m and w_s. The blue area shows the efficiency unit pairs (a_1, a_2) where L is the optimal sector, the red area shows where M is optimal, and the green area shows where S is optimal.

income earned by all household members is:

$$m = w_l L_l + w_m N_m + w_s N_s \tag{16}$$

The household has to allocate the total income across the three goods, and assign the home production that each member has to provide. We assume that the household maximizes the following utility:

$$\begin{aligned}
\max_{c_{lh}, c_{lm}, c_m, c_s, h} \; & u(c_{lh} + c_{lm}, c_m, c_s) - v(h) \\
& \ln\left(\theta_l (c_{lm} + c_{lh} - \gamma_l)^{\frac{\varepsilon-1}{\varepsilon}} + \theta_m c_m^{\frac{\varepsilon-1}{\varepsilon}} + \theta_s (c_s + \gamma_s)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}} - \frac{\phi}{\nu}\left(\frac{c_{lh}}{A_l}\right)^{\nu} \\
\text{s.t.} \quad & p_l c_{lm} + p_m c_m + p_s c_s \leq m \\
& c_{lh} = h A_l \\
& 0 \leq c_{lh}, clm, c_m, c_s
\end{aligned}$$

where m is given by (16), p_l, p_m, and p_s are the prices of the low-skilled services, the manufacturing goods, and the high-skilled services. Note that we assume that low-skilled services produced at home, c_{lh}, and bought on the market, c_{lm}, are perfect substitutes We assume that the disutility from home production hours takes the following form:

$$v(h) = \frac{\phi}{\nu} h^{\nu}.$$

The hours spent on home production cause disutility to the individual, which is increasing ($\phi > 0$) in the number of hours spent on home production, and potentially convex ($\nu > 1$). The number of hours needed per household member to attain c_{lh} units of low-skilled services is $h = c_{lh}/A_l$, thus as the productivity in the low-skilled sector increases, the number of hours needed in home production to achieve the same output fall.[7]

The utility of the consumer is non-homothetic in low-skilled services ($c_{lh} + c_{lm}$), manufacturing goods (c_m), and high-skilled service goods (c_s):

$$u(c_{lh} + c_{lm}, c_m, c_s) = \ln\left(\theta_l(c_{lm} + c_{lh} - \gamma_l)^{\frac{\varepsilon-1}{\varepsilon}} + \theta_m c_m^{\frac{\varepsilon-1}{\varepsilon}} + \theta_s(c_s + \gamma_s)^{\frac{\varepsilon-1}{\varepsilon}}\right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

There is a subsistence amount of low-skilled services the household needs, $\gamma_l > 0$, thus $c_{lh} + c_{lm}$ have to exceed this level. We assume that $\gamma_s > 0$, which is equivalent to assuming that c_s is a luxury good, i.e. the household only demands a positive amount if it is sufficiently rich.
The following corollary summarizes the solution of the household's problem.

Corollary 1. *The following cases can arise:*

1. *At the lowest income levels ($m < m_{12} \equiv \min\{m_{12a}, m_{12b}\}$) households spend all their income on manufacturing goods, and home produce low-skilled services. In such an income range the optimal choices are:*

$$\begin{aligned} c_s &= c_{lm} = 0 \\ c_m &= \frac{m}{p_m} \\ c_{lh} \;\; &\textit{solves} \;\; v'(c_{lh}) = u_1\left(c_{lh}, \frac{m}{p_m}, 0\right) \end{aligned}$$

2. *For medium income levels ($m \in [m_{12}, m_{23})$, where $m_{23} \equiv \min\{m_{2a3}, m_{2b3}\}$) households either start to spend on low- or high-skilled services,*

   A. *if the household spends on S goods ($m \in [m_{12a}, m_{2a3})$ and $m_{12a} < m_{12b}$), then*

$$c_{lm} = 0$$

   and the rest jointly solve

$$\begin{aligned} \frac{u_2(c_{lh}, c_m, c_s)}{u_3(c_{lh}, c_m, c_s)} &= \frac{p_m}{p_s} \\ u_1(c_{lh}, c_m, c_s) &= v'(c_{lh}) \\ p_m c_m + p_s c_s &= m \end{aligned}$$

[7] Given these assumptions on the disutility of home production hours, household members always do a positive amount of home production. Naturally there should be a maximum limit on home production hours, but we assume that the preference and productivity parameters are such that this limit never becomes binding.

B. if the household spends on low-skilled services ($m \in [m_{12b}, m_{2b3})$ and $m_{12b} < m_{12a}$), then

$$c_s = 0$$

and the rest jointly solve

$$\frac{u_2(c_{lh} + c_{lm}, c_m, 0)}{u_1(c_{lh} + c_{lm}, c_m, 0)} = \frac{p_m}{p_l}$$
$$u_1\left(c_{lh} + c_{lm}, c_m, 0\right) = v'(c_{lh})$$
$$p_l c_{lm} + p_m c_m = m$$

3. *For the highest income levels, if $m > m_{23}$ the optimal choices jointly solve*

$$\frac{u_2(c_{lh} + c_{lm}, c_m, c_s)}{u_1(c_{lh} + c_{lm}, c_m, c_s)} = \frac{p_m}{p_l}$$
$$\frac{u_2(c_{lh} + c_{lm}, c_m, c_s)}{u_3(c_{lh} + c_{lm}, c_m, c_s)} = \frac{p_m}{p_s}$$
$$u_1\left(c_{lh} + c_{lm}, c_m, c_s\right) = v'(c_{lh})$$
$$p_l c_{lm} + p_m c_m + p_s c_s = m$$

The cutoff m_{12a} is the income level, where the demand for S goods is exactly zero, $c_s = 0$, under case 2a. The cutoff m_{2a3} is defined as the income level where the demand for c_{lm} is exactly zero under case 3 (therefore it can only follow case 2a). The cutoff m_{12b} is the income level, where the demand for c_{lm} is exactly zero according to case 2b. The cutoff m_{2b3} is the income level where the demand for c_s is exactly zero under case 3 (therefore it can only follow case 2b).

Proof. See Appendix □

3.3 Competitive Equilibrium

A competitive equilibrium is given by cutoff sector of work abilities $\{a_{1lm}, a_{2ls}, a_{2ms}(a_1)\}$, wages $\{w_l, w_m, w_s\}$, prices $\{p_l, p_m, p_s\}$, and consumption demands $\{c_{lm}, c_m, c_s\}$, given productivities $\{A_l, A_m, A_s\}$ which satisfy:

1. $\{a_{1lm}, a_{2ls}, a_{2ms}(a_1)\}$ are given by (7), (8), and (9);
2. sector of work decisions are made according to (10), (11), and (12);
3. the unit wage rates are such that the market for L, M and S labor clears, i.e. (6), (4), and (2) are satisfied;

4. p_l, p_m and p_s are such that the market for L, M and S goods clears, i.e.

$$c_{lm} = A_l L_l \tag{17}$$

$$c_m = A_m N_m \tag{18}$$

$$c_s = A_s N_s \tag{19}$$

where c_{lm}, c_m, c_s are given by Corollary 1, L_l, N_m, N_s are given by (13), (14), and (15).

In the next section we analyze the behavior of the model as productivity levels in manufacturing and high-skilled services increase.

4 Quantitative results

In this section we quantitatively assess the contribution of structural transformation to the polarization of employment and wages across industries. To do this we consider the evolution of the competitive equilibrium in terms of employment shares and relative average sectoral wages as productivity increases in manufacturing and in high-skilled services. We first describe the data targets and the calibration strategy, and then discuss the quantitative importance of our mechanism.

4.1 Data targets

We calibrate our model to replicate four key moments of the US economy in 1950 and in 1960. These moments are the relative average industry wages, and the industry employment shares. Data for the average industry wages and the industry employment shares come from the 1950 and 1960 US Census data. Each employed individual is categorized as a high-skilled service, manufacturing or low-skilled service worker, based on their industry code (ind1990). Employment shares are calculated as share of hours worked, as in section 2.

4.2 Calibration

All parameters are time-invariant, and the only exogenous change over time is productivity growth. Herrendorf, Rogerson, and Valentinyi (2013) estimate a labor-augmenting technological progress on value added output, with annual growth rates of 2% for manufacturing, and 1.1% for the service sector, which we assume to be the growth rate of the high-skilled services. We assume that there is no productivity growth in low-skilled services.

The following parameters need to be calibrated: the initial productivity levels A_l, A_m, A_s, the parameters of the utility function $\theta_l, \theta_m, \theta_s, \varepsilon, \gamma_l, \gamma_s, \nu, \phi$, and the distribution of abilities, $f(a_1, a_2)$.

We need to identify the distribution of innate abilities, $f(a_1, a_2)$. We assume that the abilities are distributed uniformly, and that the ability in manufacturing is independent of the ability in high-skilled

services. We normalize the mean of a_1 and a_2 to be unity, as these means cannot be separately identified from other parameters of the model. Given these assumptions on the distribution, two parameters are left to be calibrated, the minimum (and thus the maximum, given that the mean is one) of both abilities: $\underline{a}_1$ and $\underline{a}_2$. Given the assumptions on the distribution of abilities, in the context of our model any observed sectoral raw labor shares uniquely identify the sector of work cutoffs. Given the sector of work cutoffs, we can calculate the average relative wages they imply. To see this mathematically consider similarly to the effective labor supplied to each sector, the raw labor shares:

$$
\begin{aligned}
L_l &= \int_{\underline{a}_1}^{a_{1lm}} \int_{\underline{a}_2}^{a_{2ls}} f(a_1, a_2) da_2 da_1, \\
L_m &= \int_{a_{1lm}}^{\overline{a}_1} \int_{\underline{a}_2}^{a_{2ms} a_1} f(a_1, a_2) da_2 da_1, \\
L_s &= \int_{\underline{a}_1}^{\overline{a}_1} \int_{\max(a_{2ls}, a_{2ms} a_1)}^{\overline{a}_2} f(a_1, a_2) da_2 da_1.
\end{aligned}
$$

Matching the raw employment shares for any ability distribution gives us the cutoffs a_{2ls}, a_{1lm}, and a_{2ms}. Given these cutoffs the relative average wages implied by the model can be calculated as:

$$
\begin{aligned}
\frac{\overline{w}_l}{\overline{w}_m} &= \frac{w_l}{w_m \frac{N_m}{L_m}} = a_{2ls} \frac{L_m}{\int_{a_{1lm}}^{\overline{a}_1} \int_{\underline{a}_2}^{a_{2ms} a_1} a_1 f(a_1, a_2) da_2 da_1}, \\
\frac{\overline{w}_s}{\overline{w}_m} &= \frac{w_s \frac{N_s}{L_s}}{w_m \frac{N_m}{L_m}} = \frac{1}{a_{2ms}} \frac{L_m \int_{\underline{a}_1}^{\overline{a}_1} \int_{\max(a_{2ls}, a_{2ms} a_1)}^{\overline{a}_2} a_2 f(a_1, a_2) da_2 da_1}{L_s \int_{a_{1lm}}^{\overline{a}_1} \int_{\underline{a}_2}^{a_{2ms} a_1} a_1 f(a_1, a_2) da_2 da_1},
\end{aligned}
$$

where we used that the cutoffs in the model are ratios of the unit wage rates: $a_{2ls} = w_l / w_s$, and $a_{2ms} = w_m / w_s$.

Thus given a distribution $f(a_1, a_2)$, matching raw labor shares implies relative average wages in the context of our model. Therefore we can pin down $\underline{a}_1, \underline{a}_2$, to be such that when matching the raw employment shares in 1950, the model also matches the relative average wages in 1950.

We normalize all initial productivity values to 1, i.e. $A_l(0) = A_m(0) = A_s(0) = 1$. In the literature, when sectoral output is measured in value added terms, a very low value of ε is estimated and used (Herrendorf, Rogerson, and Valentinyi (2013)). In the calibration we test the sensitivity of our results to this parameter, but in the baseline case we set it to 0.02. Without loss of generality we assume that the weights on the three types of goods in the utility function sum to one: $\theta_l + \theta_m + \theta_s = 1$. We are thus left to calibrate 6 parameters of the utility function: $\theta_l, \theta_m, \gamma_l, \gamma_s, \nu, \phi$. We calibrate these parameters by minimizing the squared deviation of the following model and data moments: raw employment shares in 1950, relative average wages in 1950 and 1960. The summary of the calibrated parameters are in the Table 2.

Table 2: Calibrated Parameters

	Description	Value
A_l	1950 productivity in L	1
A_m	1950 productivity in M	1
A_s	1950 productivity in S	1
γ_L	subsistence amount of home good	0.2351
γ_S	degree of luxury of c_s	0.3551
ε	CES b/w L, M and S in consumption	0.02
θ_L	weight on $c_{lm} + c_{lh}$	0.1806
θ_M	weight on c_m	0.2575
θ_S	weight on c_s	$\theta_S = 1 - \theta_M - \theta_L$
ϕ	relative weight on disutilty from h	0.3631
ν	curvature of disutility from h	1.0755
$[\underline{a}_1, \overline{a}_1]$	range of manufacturing efficiency	$[0.4560, 1.5440]$
$[\underline{a}_2, \overline{a}_2]$	range of high-skilled service efficiency	$[0, 2]$

4.3 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 study the endogenous response of employment and wages. Our baseline is, as in the data, that productivity growth is higher in manufacturing than in services. Taking estimates by Herrendorf, Rogerson, and Valentinyi (2013) we let M' s productivity grow by 2% per annum and S's productivity grow by 1.1%.

Figure 6 plots the resulting dynamics. Productivity growth in M and S implies that M and S goods become relatively cheaper. Since S is a luxury good, the demand for it increases more than the demand for M. This implies that more labor needs to be employed in sector S, which is exacerbated by the higher productivity growth in M than in S. To satisfy demand, employment in the S sector has to increase, and firms have to pay higher unit wages to attract workers. The highest S-ability workers who would have worked in manufacturing before, now work in sector S. S sector unit wages have to increase to attract a sufficient amount of new entrants. As a consequence, the cutoff ability of workers sorting into S falls, which tends to decrease the average wage paid in the sector. However, the effect of higher unit wages is the dominating one, and on average S workers' wages increase relative to M. This is reflected in the gradual increase of the employment share in sector S, and the rise in the relative S wages, as can be seen in the two right panels in the bottom row of Figure 6. Hence, both wages and employment at the top-end of the distribution increase.

Moreover, with national income the demand for low-end services increases, and as households get richer, they meet a smaller fraction of this increased demand by home production, hence employment of L labor increases. As the demand for L services increase, also the wage for L workers rises, in order to pull workers into the low-skilled services sector. As a consequence, employment and wages rise at both ends of the distribution.

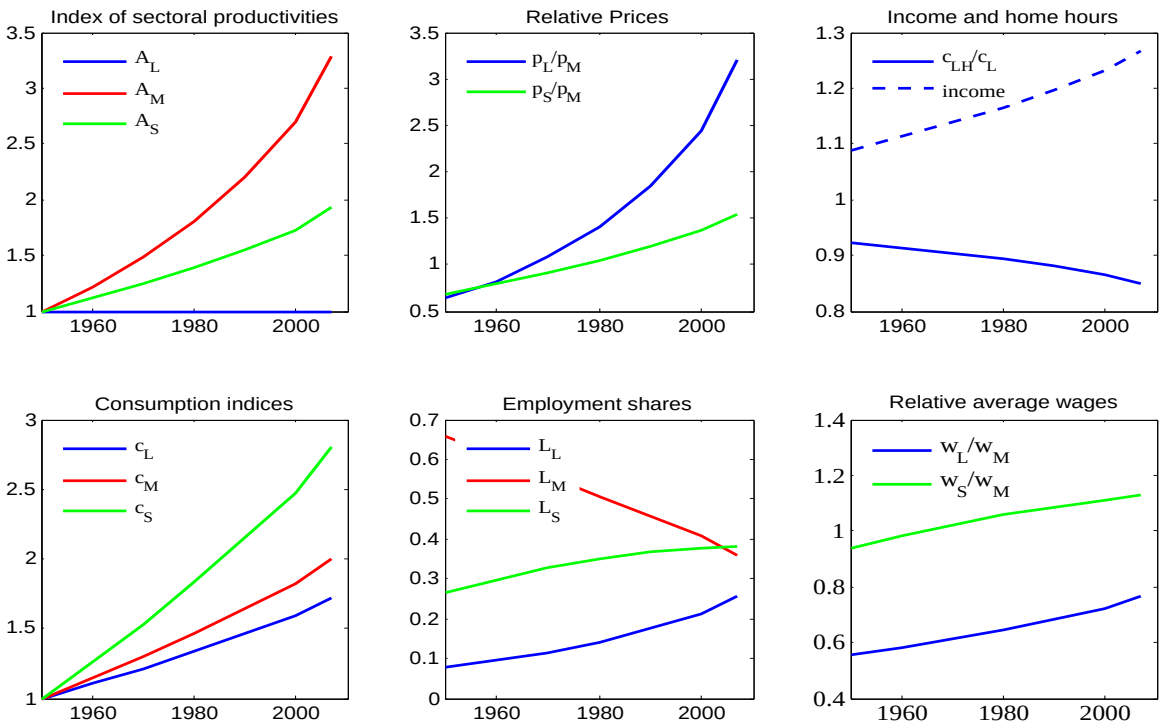

Figure 6: Transition of the benchmark model

5 Conclusions

The literature on polarization of employment and wages has typically focused on occupations. We present a set of new empirical facts that suggest that in addition to reallocations between occupations within industries, also shifts between industries contribute to the polarization of labor markets. Moreover, we show that in terms of broadly defined industries, polarization was present as early as 1950-1960 and directly linked to the decline of manufacturing employment. Based on this evidence we propose a novel explanation, one based on structural change. We develop a multi-sector model with the most parsimonious setup that yet allows heterogeneity in wages, and in which on the consumption side high-end services are luxury goods and low-end services a substitute for home production, to show that structural change over time increases employment and wages in both the low-skilled and the high-skilled service sector, thus leading to the polarization of the labor market.

References

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

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

AUTOR, D. H., AND D. DORN (2012): "The Growth of Low Skill Service Jobs and the Polarization of the U.S. Labor Market," *The American Economic Review*.

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

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

BOPPART, T. (2011): "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences," Working paper no. 2, University of Zurich Department of Economics.

BUERA, F. J., AND J. P. KABOSKI (2009): "Can Traditional Theories of Structural Change Fit the Data?," *Journal of the European Economic Association*, 7(2-3), 469–477.

——— (2012a): "The Rise of the Service Economy," *American Economic Review*, 102(6), 2540–69.

——— (2012b): "Scale and the origins of structural change," *Journal of Economic Theory*, 147(2), 684 – 712, ¡ce:title¿Issue in honor of David Cass¡/ce:title¿.

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

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

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

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

——— (2011): "Explaining job polarization: the roles of technology, offshoring and institutions," Open Access publications from Katholieke Universiteit Leuven urn:hdl:123456789/331184, Katholieke Universiteit Leuven.

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

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

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

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

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

MAZZOLARI, F., AND G. RAGUSA (2007): "Spillovers from high-skill consumption to low-skill labor markets," *The Review of Economics and Statistics*, (0).

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

MICHAELS, G., A. NATRAJ, AND J. VAN REENEN (2010): "Has ICT Polarized Skill Demand? Evidence from Eleven Countries over 25 years," Working Paper 16138, National Bureau of Economic Research.

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

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

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

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 (2012) we restrict the sample to individuals who were in the labor force and of age 16 to 64 in the year preceding the survey. We drop residents of institutional group quarters and unpaid family workers. We also drop respondents with missing earnings or hours worked data and those who work in agricultural occupations or industries or in the military. Our employment measure is the product of weeks worked times usual number of hours per week. We also compute hourly wages as earnings divided by the product of usual hours and weeks worked.[8]

To construct the 30-year change graphs of figure 1 and 10-year change graphs of figure **??** we follow the methodology used in Autor, Katz, and Kearney (2006), Acemoglu and Autor (2011), and Autor and Dorn (2012), which requires to have a balanced panel of occupations. Dorn (2009) and Autor and Dorn (2012) provide a balanced panel of occupational classifications ('*occ1990dd*') over 1980-2008, which we take to construct a balanced panel over 1950-2007 by aggregating occupational codes as needed. This leaves us with 183 balanced occupational codes.

Following the routinization literature, we then classify occupations into three categories:

1. 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, health-care support

2. Routine:
   construction trades, extractive, machine operators, assemblers, inspectors, mechanics and repairers, precision production, transportation and material moving occupations
   sales, administrative support

3. Abstract – skilled non-routine:
   managers, management related, professional specialty, technicians and related support

Based on our theory we classify the industries into three sectors:

1. Low-skilled services – substitutes for household production:
   retail bakeries, eating and drinking places, personal services, entertainment, business and repair services (except advertising and computer and data processing services), nursing and personal

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

Figure 7: Wage and employment polarization II.

??

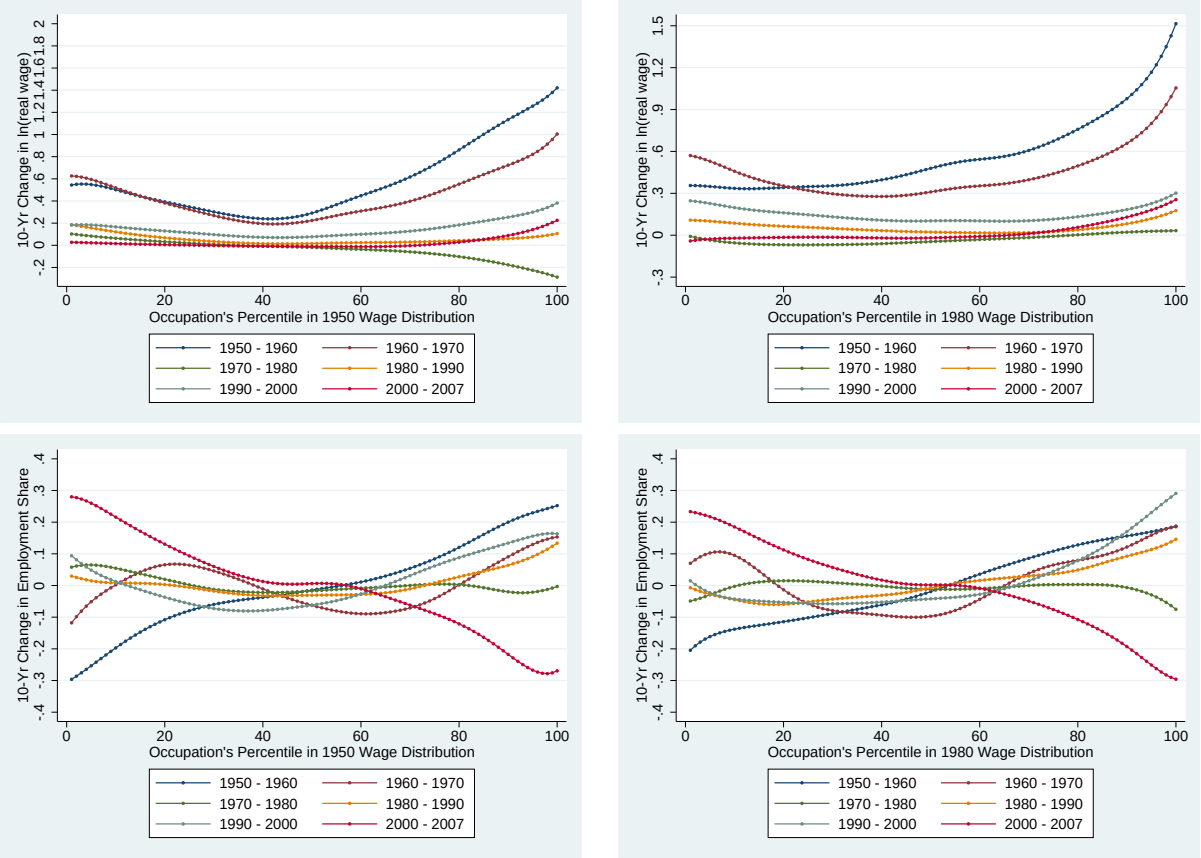

Notes: Wages are calculated from IPUMS US Census data of 1950, 1960, 1970, 1980, 1990, 2000 and American Community Survey (ACS) of 2007. Balanced occupation categories (185 of them) were defined by the authors based on Meyer and Osborne (2005) and Autor and Dorn (2012). The bottom two panels show the 10-year change in employment shares (calculated as hours supplied), and the top two panels show the 10-year change in log hourly real wages. In the left panels occupations are ranked based on their 1950 average wage (again labor supply weighted), whereas in the right panels they are ranked according to their 1980 average wage.

care facilities, child day care service, family child care homes, residential care facilities, social services, detective and protective services

2. Manufacturing:
   mining, construction, manufacturing, transport and public utilities, wholesale trade, retail trade

3. High-skilled services – luxury goods
   communications, transportation, finance, insurance and real estate, theaters and motion pictures, professional and related services, public administration, advertising, computer and data processing services

Table 3: Descriptive statistics by industry

	low-skilled services	manufacturing	high-skilled services
Highschool Dropout	23.79%	23.64%	7.41%
Highschool Graduate	34.95%	38.02%	22.81%
Some College	27.90%	24.47%	28.86%
College Degree	10.35%	10.94%	24.07%
Postgraduate	3.01%	2.93%	16.85%
Mean Years of Education	12.21	12.19	14.21
Female Share	54.12%	27.59%	54.72%
Foreign-Born Share	16.18%	10.28%	8.97%

Figure 8: Counterfactual exercise: only changes in the gender-age composition of the labor force

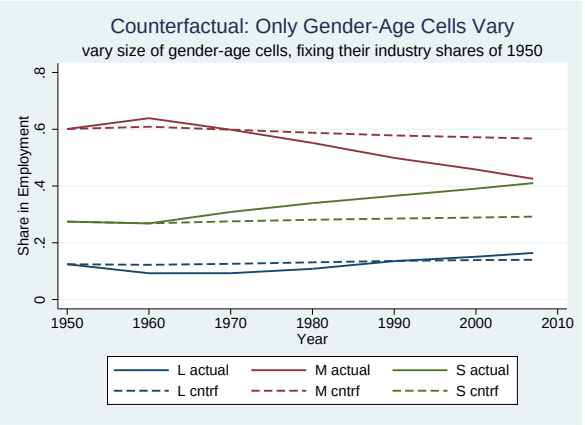

Notes: Employment shares (in terms of hours) are calculated from the same data as in Figure 2 and Table 1. The actual data is shown as solid lines, while the dashed line show how the employment shares of industries would have evolved if only the relative size of gender-age cells in the labor force had changed over time.