

Job polarization, structural transformation and biased technological change ^{*}

Zsófia L. Bárány Christian Siegel

October 2018

Abstract

By reviewing our work in Bárány and Siegel (2018a, 2018b), this article emphasizes the link between job polarization and structural change. We summarize evidence that job polarization in the United States 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. Furthermore, 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, and these two phenomena are closely linked. Finally, we propose a model where technology evolves at the sector-occupation cell level that can capture the employment reallocation across sectors, occupations, and within sectors. We show that this framework can be used to assess what type of biased technological change is the driver of the observed reallocations. The data suggests that technological change has been biased not only across occupations or sectors, but also across sector-occupation cells.

Keywords: biased technological change, structural change, employment polarization
JEL codes: O41, O33, J24

^{*} This article reviews findings of our previous joint work, and was prepared for the conference "Polarization(s) in Labor Markets" organized by DARES and the ILO in Paris on June 19, 2018. Zsófia Bárány: Sciences Po and CEPR. Email: zsofia.barany@sciencespo.fr, Christian Siegel: University of Kent, School of Economics and Macroeconomics, Growth and History Centre. Email: c.siegel@kent.ac.uk

1 Introduction

Over the last several decades the labor markets in most developed countries have experienced substantial changes. Since the middle of the twentieth century there has been structural change, the movement of labor out of manufacturing and into the service sectors. One of the key explanations for structural transformation is differential productivity growth – or biased technological progress – across sectors, combined with complementarity between the goods and services produced by different sectors (Ngai and Pissarides (2007)).¹ At the level of occupations several papers documented the polarization of labor markets in the United States and in several European countries since the 1980s: employment shifted out of middle-earning routine jobs to low-earning manual and high-earning abstract jobs. The main explanation for this phenomenon is the routinization hypothesis, which assumes that information and computer technologies (ICT) substitute for middle-skill, routine occupations, while they complement high-skill, abstract occupations; in other words technological progress that is biased across occupations (Autor, Levy, and Murnane (2003), Autor, Katz, and Kearney (2006), Autor and Dorn (2013), Goos, Manning, and Salomons (2014)). Both literatures study the impact of differential productivity growth. One focuses on the productivity across sectors and its interaction with the demand for goods and services, while the other focuses on the productivity of tasks or occupations, and its impact on the relative demand for these occupations. In this paper we review our previous work which suggests that these two phenomena are connected and should not be studied in isolation, especially in order to understand the driving forces behind the reallocation of labor across sectors and occupations.

In Bárány and Siegel (2018a) we show that polarization started much earlier than previously thought, and that it is closely linked to the structural transformation of the economy. This on its own suggests that there might be a common driving force behind structural transformation and polarization. In Bárány and Siegel (2018b) we go further; we demonstrate that there is an even tighter connection between the sectoral and occupational reallocation of employment, and we explicitly study the technological changes underlying both.

In Bárány and Siegel (2018a) we document first that in the US occupational polarization both in terms of wages and employment has started in the 1950s, much earlier than suggested by previous literature. Second, we show that a similar polarization pattern is present for broadly defined sectors of the economy, low-skilled services, manufacturing, and high-skilled services. Moreover, we show that a significant part of the occupational employment share changes is driven by shifts of employment across sectors, and that sectoral effects also explain a large part of occupational wage changes. These findings suggest that the decline in routine employment is strongly connected to the decline in manufacturing employment. We propose a model to show that differences in productivity growth across sectors lead

¹ Some papers emphasize changes in the supply of an input which is used at different intensity across sectors (Caselli and Coleman (2001), Acemoglu and Guerrieri (2008)). Other papers study the role of non-homothetic preferences, where changes in aggregate income induce a reallocation of employment across sectors (Kongsamut, Rebelo, and Xie (2001), Boppart (2014)).

to the polarization of wages and employment at the sectoral level, which in turn imply polarization in occupational outcomes.

In Bárány and Siegel (2018b) we look at the data from a different perspective: we study employment patterns across sector-occupation cells in the economy. We document some trends in occupation and sector employment that have not received much attention in the literature. First, the manufacturing sector has the highest share of routine workers; by far most of the decline in routine employment occurred in manufacturing, and conversely almost all of the contraction in manufacturing employment occurred through a reduction in routine employment. Second, the high-skilled service sector has the highest share of abstract workers; most of the expansion in abstract employment happened in the high-skilled service sector, and most of the increase in high-skilled service employment was due to an expansion in abstract employment. These patterns reveal that the sectoral and the occupational reallocation of employment are closely linked. Furthermore, the overlap of occupations and sectors implies that it is hard to identify the technological changes which underlie the observed labor market patterns. To overcome this issue, we specify a flexible model of the production side of the economy in which technological change can be biased towards workers in specific sector-occupation cells. We use key equations of this model together with data from the US Census and from the U.S. Bureau of Economic Analysis to draw conclusions about the bias in productivity changes across sector-occupation cells.

This approach departs from the recent literature connecting the phenomena of structural change and polarization across occupations in that we do not a priori restrict the nature of technological change. Goos et al. (2014) suggest that differential occupation intensity across sectors and differential occupational productivity growth can lead to employment reallocation across sectors. Duernecker and Herrendorf (2016) show in a two-sector two-occupation model that unbalanced occupational productivity growth by itself provides dynamics consistent with structural change and with the trends in occupational employment, both overall and within sectors. Lee and Shin (2017) allow for occupation-specific productivity growth and find that their calibrated model can quantitatively account for polarization as well as for structural change, and in an extension find a limited role for sector-specific technological change. Aum, Lee, and Shin (2018) analyze the role of routinization (differential productivity growth of occupations) and computerization across industries as well as industry-specific TFP differences in the recent productivity slowdown, and find in their model with homogeneous labor that sectoral TFP differences have a rather small effect.

The close link in the data between the sectoral and occupational reallocation of labor explains why models which allow for productivity growth differences only at the sectoral or only at the occupational level can go a long way in accounting for the reallocations across both dimensions. However, such restricted models load all differences in technological change on one type of factor, therefore not allowing to identify whether these differences arise indeed at the level of sectors or of occupations. We view our framework as an important and useful first step in identifying the true bias in technological change. In this article we explain how certain aspects of the data can be used to draw qualitative conclusions, whereas in Bárány and Siegel (2018b) we use a richer methodology to quantify the bias in

technology across sector-occupation cells and to decompose it further into common components. To summarize our results, we find that technological change has been biased in more nuanced ways, not just across occupations or sectors, but across sector-occupation cells.

2 A historical perspective on polarization

In Bárány and Siegel (2018a) we use data from the US Census between 1950 and 2000 and the 2007 ACS to study the patterns of employment and wages both across occupations and across sectors. In the following three subsections we summarize the main empirical results we established there. Our main findings are the following: (1) occupational polarization both in terms of wages and employment started as early as 1950 in the US, (2) wage and employment polarization is also visible in terms of broadly defined industries, (3) a large part of polarization in terms of occupations is driven by changes at the level of industries. In subsection 2.4 we go further and document the changes in employment at the sector-occupation cell level where we see a strong overlap between the evolution of occupational and sectoral employment trends.

2.1 Occupational polarization

Figure 1 plots the smoothed changes in log real wages and employment shares for occupational percentiles, with occupations ranked according to their 1980 mean hourly wage, following the methodology used in Autor et al. (2006), Acemoglu and Autor (2011), and Autor and Dorn (2013). Departing from the literature, we do not restrict attention to recent years but show the changes starting from 1950 for different 30-year periods. The left panel shows that there has been (real) wage polarization throughout, as occupations towards the middle of the wage distribution gained less than occupations at both extremes. The right panel shows that also in terms of their shares in hours worked, middle earning occupations have been tending to do worse than both low- and high-earning occupations. Though the pattern is less striking than for wages, polarization of employment has occurred since the 1950s.

To get a sense of which occupations are driving these changes and whether there are any significant differences across decades, in Figure 2 we show for 10 coarser occupational categories the decade-by-decade change in total hours worked and mean log wages. The categories we use follow Acemoglu and Autor (2011), and are ranked according to the occupations' mean wages, from lowest earners on the left to highest earners on the right. Between 1950 and 1960 a clear pattern cannot be discerned, whereas from 1960 onwards, it is clear that both total hours worked and mean log wages grew faster at both extremes than for occupations in the middle.

Finally, following Acemoglu and Autor (2011), we classify occupations into manual, routine and abstract categories.² Figure 3 plots their paths of relative wages and of employment shares. The left panel shows the path of occupational premia. These premia are the exponents of the coefficients on occupation dummies, obtained from a regression of log wages controlling for gender, race, a polynomial in potential experience, as well as occupation dummies. Obtaining the occupation premia from

² See box in section 2.3 for details of which 1-digit occupational codes are in each.

these regressions allows us to disregard changes in wage differences across occupations which are potentially caused by age, gender, or racial composition differences. It is worth to note that, as expected, the manual premium is less than the routine, while the abstract premium is the largest. 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. 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 compression until 1960, has been steadily increasing. Thus, the middle earning group, the routine workers, lost both in terms of relative average wages and in terms of the employment share to the benefit of manual and abstract workers.

All these figures constitute evidence that at the occupational level there has been employment and wage polarization in the US since at least the 1960s.

2.2 Sectoral polarization

Similar patterns can be discerned when considering the economy in terms of three broad sectors, low-skilled services, manufacturing, and high-skilled services. As common in the structural change literature our manufacturing category includes mining and construction (e.g. as in Herrendorf, Rogerson, and Valentinyi (2013)), and we split services in two (e.g. as in Buera and Kaboski (2012), Duarte and Restuccia (2017), Duernecker, Herrendorf, and Valentinyi (2017)). Classification of economic activities into broad sectors for the purpose of a model should be such that industries within sectors are very good substitutes, while they are complements across sectors. Since the service sector as a whole includes very different types of services, by splitting it in two, we improve with regards to this criterion.³

Figure 4 plots for these three sectors how wage premia and shares of hours worked evolved over time. Similarly to the occupational premia, these sector premia are calculated from a Mincerian log wage regression as the exponents of the coefficients on sector dummies, where we also control for gender, race, and a polynomial in potential experience. By construction, these sector premia do not contain changes in wage differences across sectors which are potentially caused by age, gender, or racial composition differences. As the left panel of the figure shows workers in low-skilled services typically earn less and workers in high-skilled services more per hour than those in the manufacturing sector. Moreover, it reveals that there has been a pattern of wage polarization in terms of sectors, as the wage premia in low- and in high-skilled services have been increasing since the 1960s relative to manufacturing. The right panel of the figure shows the evolution of employment shares across sectors. Manufacturing employment has been falling since the 1960s, while employment in both low- and high-skilled services has been increasing. Putting it differently, there has been employment polarization at the sectoral level as the employment share of the middle-earning sector has declined relative to both the low- and high-end sectors.

³ See boc in section 2.3 for details of which industries are in each sector.

2.3 Quantifying the impact of sectoral changes on occupations

A standard shift-share decomposition can be used to quantify the contribution of sectoral employment share changes to each occupation's employment share changes. We denote by $\Delta E_{ot} = E_{ot} - E_{o0}$ the change in the employment share of occupation o between year 0 and t , which can be decomposed as:

$$E_{ot} = \sum_i l_{oit} \Delta E_{it} + \bar{\Delta} l_{oit} E_i,$$

where $l_{oit} = L_{iot} / L_{it}$ is the share of occupation o employment within industry i employment at time t , $E_{it} = L_{it} / L_t$ is the employment share of industry i in the economy at time t , we denote the change between period 0 and t with Δ , and with the variables without a time subscript we denote the average of the variable between period 0 and period t . The first term captures the between-industry changes, this is the change in the employment share of occupation o due to changes in the industrial composition, while the changes due to within-sector reallocations are represented by the second term.

Table 1 shows the results from this decomposition for the three broad occupational categories. We conduct this decomposition for either our 3 broad occupations and 3 broad sectors, or for 10 broad occupations and 11 broad sectors.⁴ No matter the time frame or the number of industrial/occupational categories we consider, we find that a significant part of each occupation's employment share change is driven by between-industry forces. Between 1960 and 2007 around a half of the change in the manual employment share, about a third of routine, and around a quarter of abstract employment share change is driven by changes in the industrial composition of the economy.

In a similar fashion we decompose relative occupational wage changes into a component that is due to industry effects and one that is due to occupation effects. We start from the relative average wage of an occupation compared to routine wages:

$$rw_{ot} \circ \bar{\Delta} \frac{L_{iot} w_{it} w_{iot}}{L_{ot} w_{rt} w_{it}},$$

where L_{iot} / L_{ot} denotes the fraction of occupation o workers employed in industry i in period t , w_{it} / w_{rt} denotes the ratio of the average wage in industry i relative to the average wage of routine occupations in period t , and w_{iot} / w_{it} denotes the wage premium of occupation o in industry i in period t . We implement the three-way decomposition as follows. The *occupation effect* is the change in the occupational

⁴ These 11 categories 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.

wage premium within each industry relative to the industry average (w_{iot} / w_{it}). The *industry effect* is made of two parts: workers within an occupation move across industries which have different wages (L_{iot} / L_{ot}), and from changes in each industry’s average wage compared to routine wages (w_{it} / w_{rt}). Table 2 shows the results of this decomposition. It is apparent in this table that both manual and abstract occupations have been gaining in terms of wages relative to routine workers. Furthermore this table shows that more than half of occupational wage changes can be due to industry effects: due to either the reallocation of manual or abstract workers to industries with higher wages, or by faster wage growth in those industries where manual or abstract workers are employed more intensively.

2.4 Overlap between occupational and sectoral employment

While the shift between sectors per se has implications for occupational outcomes, it is informative to consider the evolution of employment at the level of sector-occupation cells since there are several distinct patterns. For the three broad sectors and the three occupational categories defined above, Figure 5 plots the evolution of sector-occupation employment shares in the U.S. between 1960–2007. The black lines show the employment share of each sector (manufacturing, low- and high-skilled services), which is then broken down into manual, routine, and abstract occupations. The economy’s structural transformation is apparent in the pronounced decline in the manufacturing sector’s employment and the rise in (particularly high-skilled) service sector employment. Occupational employment polarization is manifested in the fall of the share of routine occupations.

However, looking at occupations and sectors more carefully, two additional facts are apparent. First, the manufacturing sector has the highest share of routine labor. Second, by far most of the decline in routine employment occurred in manufacturing, whereas in the two service sectors it declined only slightly. Similarly, almost all of the increase in the employment share of abstract occupations took place in the high-skilled service sector, and most of the increase in manual employment up to 2000 was in low-skilled services. It is these patterns that imply that different economic models can explain both the sectoral and the occupational reallocations to a large degree through either sector- or occupation-specific technological change alone. However, as many models tend to a priori restrict attention to one form of technological bias, for instance only across sectors (as in Bárány and Siegel (2018a)) or only across occupations (e.g. as in Goos et al. (2014) or Duernecker and Herrendorf (2016)), they remain silent about the nature of the bias in technological change, despite replicating many aspects of the data.

In Bárány and Siegel (2018b) we take a different approach and propose a flexible setup that allows for productivity changes that are neutral (economy-wide), specific to firms in particular industries (producing particular products), specific to workers in certain occupations (linked to their task content), or specific to occupation-sector cells. In the next section we outline key features of this model and explain how certain aspects of the data inform us about how productivity changed differentially across sectors and occupations. One important aspect is that we focus on employment

reallocations not only between sectors and occupations, but also between occupations within sectors. Inspecting Figure 5 closely reveals for instance that routine employment declined not only overall, but also as a share within each sector. In the next section we show that observing the changes in occupational wages, within-sector shares of employment and of income, and sectoral prices, allows us to infer what type of biased technological change has been occurring.

3 Technological biases

To understand what type of technological change might be driving these phenomena, we formulate a model of the production side of the economy. There are two key assumptions in our framework. The first is that we explicitly assume that workers of different occupations are not perfect substitutes, and thus the factors of production are the labor supplied in various occupations. This formulation is based on the observation that there are significant differences in wages across occupations, and that different occupations perform different tasks. Second, we allow for different sectors to value these types of workers differently in production. In the following we outline the key features of the model and draw some conclusions about the likely biases in technological change based on the data we summarized in the previous section. In Bárány and Siegel (2018b) we go much further by providing a framework that can be used to quantify factor-augmenting technological change from objects in the data and by decomposing these further into neutral, sector, occupation, and idiosyncratic components.

3.1 Assumptions: The production side of the economy

The three sectors in the economy respectively produce in perfect competition low-skilled services (L), manufacturing (M), and high-skilled services (H). Labor is the only input in production, but differentiated in terms of occupations. Each sector $J \in \{L, M, H\}$ employs all three types of occupations (manual, routine and abstract), with the following CES production function:

$$Y_J = \left[\left(a_{mJ} l_{mJ} \right)^{\frac{h-1}{h}} + \left(a_{rJ} l_{rJ} \right)^{\frac{h-1}{h}} + \left(a_{aJ} l_{aJ} \right)^{\frac{h-1}{h}} \right]^{\frac{h}{h-1}} \quad (1)$$

where $\eta \in [0, \infty)$ is the elasticity of substitution between the different types of labor, l_{oJ} is occupation o labor used in sector J , and $a_{oJ} > 0$ is an *sector-occupation* specific labor augmenting technology term for occupation $o \in \{m, r, a\}$ in sector J . In this formulation, a_{oJ} in the initial year reflects the initial productivity as well as the intensity at which sector J uses occupation o , whereas any subsequent change over time reflects sector-occupation specific technological change. The assumption that the productivity depends on both the sector and the occupation of the worker renders this production function very flexible, as it does not impose any restrictions on the nature of technological change. In particular, it does not require taking a stance on whether technological change is specific to sectors or occupations.

Firms in all sectors take prices and wages as given and maximize profits by choosing occupation $o \in \{m, r, a\}$ employment such that:

$$\frac{\eta p_J}{\eta l_{oJ}} = \frac{\eta \left(p_J Y_J - \sum_i w_i l_{iJ} \right)}{\partial l_{oJ}} = p_J Y_J^{\frac{1}{\eta}} a_{oJ}^{\frac{h-1}{\eta}} l_{oJ}^{\frac{-1}{\eta}} - w_o = 0. \quad (2)$$

Combining these first order conditions for different occupations, optimal relative occupational employment within sectors satisfies:

$$\frac{l_{mJ}}{l_{rJ}} = \left(\frac{w_r}{w_m} \right)^{\eta} \left(\frac{a_{mJ}}{a_{rJ}} \right)^{h-1} \quad (3)$$

$$\frac{l_{aJ}}{l_{rJ}} = \left(\frac{w_r}{w_a} \right)^{\eta} \left(\frac{a_{aJ}}{a_{rJ}} \right)^{h-1} \quad (4)$$

These expressions demonstrate how optimal relative labor demand depends on the relative wages and on the relative productivity of different occupations. *Ceteris paribus*, all sectors optimally use more manual labor relative to routine labor if the relative routine wage, w_r/w_m , is higher. Additionally, if in sector J the term $(a_{mJ}/a_{rJ})^{h-1}$ is larger then it is optimal to use relatively more manual labor in that sector. It is important to note that an improvement in the relative productivity of for example manual to routine workers, i.e. an increase in a_{mJ}/a_{rJ} , would lead to a different impact on the optimal relative labor use depending on whether η is larger or smaller than 1. If $\eta > 1$, then the different occupations are good substitutes, so the improvement in the relative productivity of manual workers would lead to an increased relative demand for manual workers. If, on the other hand $\eta < 1$ and the different workers are complements, then an improvement in relative technology would lead to a reduction in relative demand. So for example *routinization* in sector J , i.e. the replacement of routine workers by certain technologies, would be captured by an increase in $(a_{mJ}/a_{rJ})^{h-1}$ and in $(a_{aJ}/a_{rJ})^{h-1}$.

Using optimal manual and abstract labor as a function of routine labor from (3) and (4) and substituting these into (2) for routine labor, we can express sector J prices as:

$$p_J = \left[\frac{a_{mJ}^{h-1}}{w_m^{h-1}} + \frac{a_{rJ}^{h-1}}{w_r^{h-1}} + \frac{a_{aJ}^{h-1}}{w_a^{h-1}} \right]^{\frac{1}{1-\eta}}. \quad (5)$$

3.2 Inferring technological biases

The assumptions we made about the economy's production side constitute a framework which, given η , the elasticity of substitution between the different types of occupational labor within sectors, can be used to draw conclusions from the data

about the sector-occupation specific labor augmenting technologies, the α s. While there is no consensus on the exact value of η , the literature agrees that occupations tend to be complements, and therefore this elasticity of substitution has to be less than 1. Goos et al. (2014) estimate, Duernecker and Herrendorf (2016), Lee and Shin (2017) and Aum et al. (2018) calibrate the elasticity of substitution to be between 0.5 and 0.9. For this reason in what follows we assume that $\eta < 1$, that is that the different occupational labor inputs are complements in production.

Multiplying the optimality conditions (3) and (4) with w_m/w_r and w_a/w_r respectively and re-arranging we get the following expressions:

$$\frac{a_{mJ}}{a_{rJ}} = \frac{w_m}{w_r} \left(\frac{q_{mJ}}{q_{rJ}} \right)^{\frac{1}{\eta-1}}, \quad (6)$$

$$\frac{a_{aJ}}{a_{rJ}} = \frac{w_a}{w_r} \left(\frac{q_{aJ}}{q_{rJ}} \right)^{\frac{1}{\eta-1}}, \quad (7)$$

where $q_{oJ} = (w_o l_{oJ}) / (p_J Y_J)$ denotes the share of income in sector J going to occupation o workers. Note that we assume that there is perfect competition, the production function is constant returns to scale, and that the only factors of production are the different types of occupational labor, which implies that profits are zero and $\sum_o q_{oJ} = 1$. From these equations we can see that given data on relative occupational wages and on occupational income shares within sectors we can infer the evolution of relative occupational productivities within a sector.

We are primarily interested in the change in relative sector-occupation productivities within sectors over time. For this reason, in Figure 6 we plot the evolution of relative wages of different occupations relative to their 1960 values. Wages in both abstract and manual occupations increased relative to routine occupations. Overall the gain in relative wages was around 25 percent in abstract occupations and around 38 percent in manual occupations. In Figure 7 we show the evolution of relative occupational income shares in all three sectors between 1960 and 2007, relative to their 1960 values. The income share of both abstract and manual workers increased relative to routine in all three sectors albeit at a different rate. Abstract workers' income share increased the most in high-skilled services (almost 2.5 fold), in manufacturing it more than doubled, while in low-skilled services it increased by 50 percent. Manual workers' income share increased the most in manufacturing (it increased six fold), in high-skilled services it more than doubled, whereas in low-skilled services it increased but less than doubled.

It is important to note that for values of the elasticity of substitution below 1, the change in relative wages and the change in income shares imply changes of opposite sign in relative productivities. The changes in relative income shares are much larger than the changes in relative wages. The lower is η the smaller is the change implied by the change in income shares, but even for relatively low values of η it dominates the implied change coming from wages. We can therefore conclude that the

productivity of routine workers had to increase in all sectors relative to both manual and abstract workers. This is a pattern common across sectors, and it is in line with the routinization hypothesis. The relative productivity of routine workers increased, and since different occupations are complements in production in all sectors, this implies a lower relative demand for routine workers in all sectors. At the same time, the magnitude of change in relative income shares is markedly different across sectors, which point to the presence of sector-occupation specific changes in productivity.

Next we analyze the evolution of relative productivities across sectors. This is informed by the movement of relative sectoral prices. Using relative occupational productivities within sectors (equations (6) and (7)) and that $\overset{\circ}{a}q_{oJ} = 1$, we can express sectoral prices (5) in terms of observables as:

$$p_J = \frac{w_r}{a_{rJ}} \left[\left(\frac{a_{mJ} w_r}{a_{rJ} w_m} \right)^{h-1} + 1 + \left(\frac{a_{aJ} w_r}{a_{rJ} w_a} \right)^{h-1} \right]^{\frac{1}{1-h}} = \frac{w_r}{a_{rJ}} \left[\frac{q_{mJ}}{q_{rJ}} + 1 + \frac{q_{aJ}}{q_{rJ}} \right]^{\frac{1}{1-h}} = \frac{w_r}{a_{rJ}} \left[\frac{1}{q_{rJ}} \right]^{\frac{1}{1-h}}.$$

By relating prices across sectors, we can express relative sector-occupation productivities as:

$$\frac{a_{rM}}{a_{rL}} = \frac{p_L}{p_M} \left(\frac{q_{rM}}{q_{rL}} \right)^{\frac{1}{h-1}}, \quad (8)$$

$$\frac{a_{rM}}{a_{rH}} = \frac{p_H}{p_M} \left(\frac{q_{rM}}{q_{rH}} \right)^{\frac{1}{h-1}}. \quad (9)$$

These two equations show that the evolution of relative sector-occupation productivities across sectors can be inferred from changes in relative sectoral prices and in the cross-sector ratio of routine workers' income shares.

Figure 8 shows how these two objects evolved over time, compared to their 1960 values. The relative income share of routine workers in manufacturing increased by more than 30 percent relative to high-skilled services, while relative to low-skilled services it fell, by just under 10 percent. Both relative prices fluctuated a bit, but while overall there was no significant change in the relative price of low-skilled services compared to manufacturing (but it decreased slightly), the relative price of high-skilled services increased by almost 80 percent.

The trends in relative prices imply that routine workers' technology improved at a faster rate in manufacturing than in high-skilled service, and at a slightly lower rate than in low-skilled services. The changes in the relative income share of routine workers, however, point in the opposite direction. Nonetheless, unless the two just happen to offset each other, this analysis highlights that routine workers' productivity changed not in the same way across sectors. For the range of the elasticity of substitution considered in the literature, i.e. $\eta \in (0.5, 0.9)$, stronger conclusions can be drawn. Given the documented data, the implied change coming from income shares

dominates, implying that routine workers' productivity in manufacturing grew faster than in low-skilled services, but it grew slower than in high-skilled services.

More generally, interpreting the patterns in the data through the lens of our model suggests that technological change has been biased across sector-occupation cells – a pure bias across occupations or sectors alone is not enough to explain the data. It is of course conceivable that there are common patterns in the cell technologies, such as common occupation or sector factors, but these are not the sole drivers.

4 Conclusion

In this article we reviewed our work in Bárány and Siegel (2018a,b) on the nexus of job polarization and structural transformation, stressing the importance of biased technological change as drivers of the observed changes in labor market outcomes both at the sectoral and at the occupational level. While sectoral reallocations, which might be caused by productivity growth differences across sectors, imply for occupations changes in employment shares and in wages that are qualitatively in line with certain aspects of the data, they cannot speak to the observed within-sector changes of occupational employment shares. This suggests that technological change must have been biased in more complex ways. However, explanations of technological change affecting workers according to their occupations differentially, such as ICT technologies adversely affecting workers in routine jobs, fall short of explaining all aspects of the data as well.

We show, an occupation-bias in technology alone is not consistent with the joint observed changes in sectoral prices, occupational wages, and occupation-sector employment shares. Analyzing the data through our framework instead suggests that the productivity of routine workers relative to abstract or manual workers changed differentially across the three sectors we consider. This leaves the possibility that technological change is entirely specific to the sector-occupation cell, or that it is biased across sectors and across occupations.

References

Daron Acemoglu and David Autor (2011). 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.

Daron Acemoglu and Veronica Guerrieri (2008). Capital Deepening and Nonbalanced Economic Growth. *Journal of Political Economy*, 116(3): 467–498.

Sangmin Aum, Sang Yoon (Tim) Lee, and Yongseok Shin (2018). Computerizing industries and routinizing jobs: Explaining trends in aggregate productivity. *Journal of Monetary Economics*, 97: 1–21.

David H. Autor and David Dorn (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5): 1553–97.

David H. Autor, Frank Levy, and Richard J. Murnane (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4): 1279–1333.

David H. Autor, Lawrence F. Katz, and Melissa S. Kearney (2006). The Polarization of the U.S. Labor Market. *American Economic Review*, 96(2): 189–194.

Zsófia L. Bárány and Christian Siegel (2018a). Job Polarization and Structural Change. *American Economic Journal: Macroeconomics*, 10(1): 57–89.

Zsófia L. Bárány and Christian Siegel (2018b). Disentangling Occupation- and Sector-specific Technological Change. CEPR Discussion Paper DP12663, CEPR.

Timo Boppart (2014). Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences. *Econometrica*, 82(6): 2167–96.

Francisco J. Buera and Joseph P. Kaboski (2012). The Rise of the Service Economy. *American Economic Review*, 102(6): 2540–69.

Francesco Caselli and Wilbur John Coleman, II (2001). The U.S. Structural Transformation and Regional Convergence: A Reinterpretation. *Journal of Political Economy*, 109(3): 584–616, 2001.

David Dorn (2009). Essays on Inequality, Spatial Interaction, and the Demand for Skills. PhD thesis, University of St. Gallen. Dissertation no. 3613.

Margarida Duarte and Diego Restuccia (2017). Relative prices and sectoral productivity. Working Paper 23979, National Bureau of Economic Research.

Georg Duernecker and Berthold Herrendorf (2016). Structural Transformation of Occupation Employment. Working paper.

Georg Duernecker, Berthold Herrendorf, and Ákos Valentinyi (2017). Structural change within the service sector and the future of baumol’s disease. CEPR Discussion Paper DP12467, CEPR.

Maarten Goos, Alan Manning, and Anna Salomons (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*, 104(8): 2509–26.

Berthold Herrendorf, Richard Rogerson, and Ákos Valentinyi (2013). Two Perspectives on Preferences and Structural Transformation. *American Economic*

Review, 103(7): 2752–89.

Piyabha Kongsamut, Sergio Rebelo, and Danyang Xie (2001). Beyond Balanced Growth. *The Review of Economic Studies*, 68(4): 869–882.

Sang Yoon (Tim) Lee and Yongseok Shin (2017). Horizontal and vertical polarization: Task specific technological change in a multi-sector economy. Working Paper 23283, National Bureau of Economic Research.

Peter B. Meyer and Anastasiya M. Osborne (2005). Proposed Category System for 1960-2000 Census Occupations. Working Papers 383, U.S. Bureau of Labor Statistics.

L. Rachel Ngai and Christopher A. Pissarides (2007). Structural Change in a Multisector Model of Growth. *The American Economic Review*, 97(1): 429–443.

Box for section 2.1

Figure 1 – as is standard in the literature, e.g. Autor et al. (2006), Acemoglu and Autor (2011) and Autor and Dorn (2013) – shows smoothed changes in log real wages or in employment shares by percentiles of the occupational wage distribution, where occupations are ranked by their 'skill level', which is approximated by the average wage of workers in the given occupation in a base year. These occupations are then put into 100 bins on the horizontal axis, each representing 1 percent of employment. For such a comparison over time a balanced set of occupational codes are needed. In Bárány and Siegel (2018a) we construct the finest possible set of occupational codes that is balanced over 1950 to 2007, extending the work of Meyer and Osborne (2005) and Dorn (2009).

Box for section 2.3

Industries are classified into our three categories as follows: low-skilled services are personal services, entertainment, low-skilled transport, low-skilled business and repair services, retail trade, and wholesale trade; manufacturing also includes construction and mining; high-skilled services are professional and related services, finance, insurance and real estate, communications, high-skilled business services, utilities, high-skilled transport, and public administration. In terms of occupations, manual workers are those working in: housekeeping, cleaning, protective service, food preparation and service, building, grounds cleaning, maintenance, personal appearance, recreation and hospitality, child care workers, personal care, service, healthcare support. Routine occupations are construction trades, extractive, machine operators, assemblers, inspectors, mechanics and repairers, precision production, transportation and material moving occupations, sales, administrative support. Finally abstract occupations comprise managers, management related, professional specialty, technicians and related support workers.

Figure 1: Smoothed changes in wages and employment

Source: Bárány and Siegel (2018a). 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. Balanced occupation categories (183 of them) were defined by the authors based on meyer2005, dorn2009 and autor2013. The horizontal axis contains occupational skill percentiles based on their 1980 mean wages. 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).

Figure 2: Polarization in broad occupational categories

Source: Bárány and Siegel (2018a). These bar graphs show for ten broad occupational categories, as defined in Acemoglu and Autor (2011), the decade-by-decade percentage change in mean log wages (left panel) and in hours worked (right panel).

Figure 3: Polarization for broad occupations

Source: Bárány and Siegel (2018a). Occupational wage premia and employment shares (in terms of hours) are calculated from the same data as in Figure 1. Each worker is classified into one of three occupations based on their occupation code (for details of the occupation classification see text). The left panel shows relative wages: the abstract and the manual premium compared to routine (and their 95% confidence intervals), implied by the regression of log wages on gender, race, a polynomial in potential experience, and occupation dummies. The right panel shows employment shares, calculated in terms of hours worked.

Figure 4: Polarization for broad industries

Source: Bárány and Siegel (2018a). 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). 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 5: Sector-occupation employment

The data used is the same as in Figure 1. Each worker is classified into one of three sectors based on their industry code and one of three occupations based on their occupation code (for details of the industry and the occupation classification see text), employment shares in the entire economy are calculated in terms of hours.

Figure 6: Change in relative occupational wages

The data is taken from IPUMS US Census data for 1960, 1970, 1980, 1990, 2000 and the American Community Survey (ACS) for 2007. Each worker is classified into one of three occupations based on their occupation code (for details of the occupation classification see text).

Figure 7: Change in relative occupational income by sector

The data used is the same as in Figure 1. Each worker is classified into one of three sectors based on their industry code and one of three occupations based on their occupation code (for details of the industry and the occupation classification see text). The left panel shows the change in relative occupational income shares from 1960 in low-skilled services, the middle in manufacturing, and the right in high-skilled in services.

Figure 8: Change in relative routine income and prices across sectors

The data is taken from IPUMS US Census data for 1960, 1970, 1980, 1990, 2000 and the American Community Survey (ACS) for 2007 and the BEA. Each worker is classified into one of three sectors based on their industry code and one of three occupations based on their occupation code (for details of the industry and the occupation classification see text). The left panel shows the change in relative routine occupational income shares from 1960 across sectors, while the right panel shows the change in relative sectoral prices across sectors.

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

Source: Bárány and Siegel (2018a). 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 or 1960 to 2007. The first two columns use 3 occupations and 3 sectors, the last two use 10 occupations and 11 industries. The 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.

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	0.180	0.148	0.225	0.218
Occupation Δ	0.108	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	-0.014	-0.050	-0.077

Source: Bárány and Siegel (2018a). 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 or 1960 to 2007. The first two columns use 3 occupations and 3 sectors, columns three and four 10 occupations and 11 industries.